

Securing Cross-App Interactions in IoT Platforms

Musard Balliu
KTH Royal Institute of Technology
 Stockholm, Sweden
 musard@kth.se

Massimo Merro
University of Verona
 Verona, Italy
 massimo.merro@univr.it

Michele Pasqua
University of Verona
 Verona, Italy
 michele.pasqua@univr.it

Abstract—IoT platforms enable users to connect various smart devices and online services via reactive apps running on the cloud. These apps, often developed by third-parties, perform simple computations on data triggered by external information sources and actuate the results of computation on external information sinks. Recent research shows that unintended or malicious interactions between the different (even benign) apps of a user can cause severe security and safety risks. These works leverage program analysis techniques to build tools for unveiling unexpected interference across apps for specific use cases. Despite these initial efforts, we are still lacking a semantic framework for understanding interactions between IoT apps. The question of what security policy cross-app interference embodies remains largely unexplored.

This paper proposes a semantic framework capturing the essence of cross-app interactions in IoT platforms. The framework generalizes and connects syntactic enforcement mechanisms to bisimulation-based notions of security, thus providing a baseline for formulating soundness criteria of these enforcement mechanisms. Specifically, we present a calculus that models the behavioral semantics of a system of apps executing concurrently, and use it to define desirable semantic policies in the security and safety context of IoT apps. To demonstrate the usefulness of our framework, we define static mechanisms for enforcing cross-app security and safety, and prove them sound with respect to our semantic conditions. Finally, we leverage real-world apps to validate the practical benefits of our policy framework.

Index Terms—Cloud-based IoT platform, IoT app security, cross-app noninterference

I. INTRODUCTION

IoT platforms provide robust application support for automating the interaction and communication between Internet-connected services and smart physical devices. This interaction is enabled by simple reactive programs known as IoT apps (or applets) running on a cloud-based IoT platform, and sensing and actuating data from services and devices on behalf of a user. These apps, often developed by third-parties, are triggered by external information sources, as in “if the room temperature exceeds a threshold”, to perform actions on external information sinks, as in “open the windows”. By exposing devices such as a thermostat and a smart window to the IoT platform via, e.g., REST APIs, IoT apps can be used to implement desirable automations like “if the room temperature exceeds a threshold then open the windows”.

Driven by the appeal of end-user programming, IoT platforms such as IFTTT [24] (If This Then That), Stringify [34], and Microsoft Flow [29] support thousands of smart devices and services with millions of users running millions

of IoT apps. These platforms help users to build powerful automations by connecting IoT devices (e.g., smart homes, security cameras, and voice assistants) to online services (e.g., Google and Dropbox) and social networks (e.g., Instagram and Twitter). For instance, the IFTTT platform allows to execute IoT *applets* that include triggers, actions, and filter code. For the platform to run an applet, users need to provide their credentials to the services associated with its triggers and actions. In the previous applet that opens the window when the temperature exceeds a threshold, the user gives the applet access to the APIs for the temperature device (e.g., a Nest Thermostat [32]) and the smart window (e.g., SmartThings [35]). Additionally, applets may contain filter code for personalization, e.g., for setting the temperature threshold. If present, the filter code is invoked after a trigger has been fired and before an action is dispatched.

Recently, researchers have shown that popular IoT platforms are susceptible to attacks that may cause severe security and safety issues for the end users and the physical devices [5]. Examples of attacks include design flaws due to over privileged permission tokens [18], unexpected information leaks by seemingly harmless apps [36], and sensitive information disclosure by malicious apps [6], [11]. To protect the users against these attacks, defensive mechanisms rely on fine-grained access control and capabilities, decentralization [19] or static [7], [11] and dynamic [6] information-flow analysis.

A more subtle vulnerability concerns the unintended or malicious interaction between different apps running on behalf of the same user [13]–[15], [17], [36]. The distinctive feature of IoT apps to affect a shared physical environment such as the room temperature, may enable unintended *cross-app interactions* between IoT apps that are installed by the same user. For instance, in addition to the above-mentioned IoT app “if the room temperature exceeds a threshold then open the windows”, a user may also install the app “if I leave my work location then turn on the thermostat at home”. While the user’s intention is to use these two apps for separate purposes, the interaction between the latter and the former may open the window while the user is not at home, thus clearing a way for burglary.

Recent research identifies numerous use cases of cross-app interactions that violate specific policies, and suggests tracking dependencies across IoT apps to identify policy violations [13]–[15], [17], [36]. These mechanisms perform inter-application program analysis to track dependencies, and (man-

ual or automatic) language processing to identify semantically-related language constructs, *e.g.*, the fact that *temperature* and *thermostat* refer to related semantic constructs, despite their syntax being different. While these approaches motivate the need for analyzing security and safety risks in cross-app interactions, foundational questions related to the interaction semantics of apps, security policies, and soundness of enforcement mechanisms remain largely unexplored.

This leads us to the following research questions: (i) What is an appropriate formal model for cross-app interaction vulnerabilities? (ii) Is there a generic policy framework for security and safety that captures the essence of cross-app interactions? (iii) How do we model implicit interactions stemming from IoT-specific features like the physical environment? (iv) Can we harden enforcement mechanisms to prove soundness guarantees in our policy framework?

Contributions: To help answering these questions, we develop a process calculus for specifying and reasoning about cross-app interactions, capturing the core features of apps in IoT platforms like IFTTT and Stringify. We then propose extensional conditions to capture the essence of security and safety requirements in a system of IoT apps executing concurrently. We demonstrate the usefulness of these conditions by considering policies from real-world apps, and discuss how they can be relaxed in order to accommodate more flexible user policies. Further, we show how standard enforcement mechanisms can be adapted to check security and safety of a system of IoT apps, thus providing strong guarantees against vulnerable cross-app interactions. We think that these conditions will provide a semantic baseline for proving soundness of current and future enforcement mechanisms in the domain of IoT apps.

Our key observation is that for a system of apps to reach an unsafe configuration, a cross-app interaction should either lead to an inconsistent state that violates the intended specification for some apps, or engage in an interaction where the action of one app triggers the execution of another app. This is supported by the intuition, as well as existing real-world vulnerabilities [13]–[15], [17], [36], that an end user may consider a system of IoT apps as safe if the runtime behavior of an app in isolation is *bisimilar* to running that app in parallel with other apps in the system. Drawing on Focardi and Martinelli’s *Generalized Non Deducibility on Composition* [21], we formalize this intuition to provide a bisimulation-based characterization of *safe cross-app interaction*. Further, we provide a simple syntactic condition and prove it sound for our notion of safe cross-app interaction. We also tackle the challenge of implicit cross-app interactions and propose an extension of our semantic condition. Finally, we envision scenarios where some form of interaction across apps can be considered as safe, and show how it can be modeled in our framework via *priorities*.

Further, we focus on confidentiality and integrity policies of a system of IoT apps and propose a *termination-insensitive* bisimulation-based security condition that accommodates these policies. As standard in information-flow con-

trol [33], the condition assumes a security classification of global services and devices, and it ensures that any interference between apps respects the security classification. We propose an extension of the flow-sensitive type system by Hunt and Sands [23] for our concurrent IoT setting, and prove it sound for our security condition.

In summary, the paper provides the following contributions:

- We present a calculus for IoT apps to study security and safety in cross-app interactions. The calculus models closely the behavioral semantics of apps in IoT platforms (Section II).
- Inspired by policy requirements in real apps, we propose an extensional condition for safe cross-app interactions, as well as a syntactic condition to enforce safe interactions (Section III).
- We extend our framework to accommodate implicit app interactions and service priorities in order to tackle the challenge of false negatives and false positives, respectively (Section IV).
- We propose a flow-sensitive security types system, enforcing information-flow policies in a system of IoT apps running concurrently (Section V).

Full proofs of our results and a number of examples of IoT apps modeled in our calculus can be found in the Appendix.

II. A CALCULUS OF IOT APPS

In this section, we define our *Calculus of IoT Apps*, called *CaITApp*, to formally specify and reason about *systems of apps*, *i.e.*, sets of concurrent IoT apps running on an IoT platform, and accessing the Internet-connected services and devices of a given user. The interface between the IoT apps and the external services and physical devices, *e.g.*, Dropbox or home security camera, is defined by APIs that enable communication between the platform and the user services and devices. As common in IoT platforms like IFTTT, the platform itself maintains a *global store* with data from a user’s services and devices, which gets updated whenever there is a change in the corresponding services and devices. Each IoT app of a given user has its own *local store*, *i.e.*, local view, which may get updated whenever the execution of that app is triggered by a change in the global store.

We start the description of our calculus with some preliminary notations. We use letters $x, y, z \in \text{Service}$ to denote the IoT platform’s (global) view of a user’s services and devices. Abusing notation, we call them just services in the following. *Values*, ranged over by $v, w \in \text{Value}$ are basic values, such as booleans, integers, real numbers, strings, etc. We assume two special values: \perp and $*$. The first represents an undefined value, while the second is a placeholder that can be replaced with “any value”.

The *syntax* of our *systems* is given by the grammar:

$$\begin{aligned} \text{Sys} \ni S ::= & S \parallel S && \} \text{ parallel composition} \\ & | \text{id}[D \times P] && \} \text{ app} \end{aligned}$$

Here, $\text{id}[D \bowtie P]$ denotes an app with a *unique identifier* $\text{id} \in \mathcal{I}$, using only the global services declared in D , with the associated permissions (read and/or write), and running the process (code) P .

The syntax for service declarations is the following:

$$\begin{aligned} \text{Decl} \ni D &::= D; D && \} \text{ declaration list} \\ &| x^R && \} \text{ service to be used in read} \\ &| x^W && \} \text{ service to be used in write} \end{aligned}$$

In the following, we will write x^{RW} , as a shorthand for $x^R; x^W$.

The syntax of our *processes* for describing the code running in our IoT apps is the following:

$$\begin{aligned} \text{Proc} \ni P &::= \text{listen}(L) && \} \text{ listener} \\ &| x \leftarrow e && \} \text{ set local store} \\ &| \text{update}(x) && \} \text{ set global store} \\ &| \text{if } b \text{ then } \{P\} \text{ else } \{P\} && \} \text{ conditional} \\ &| \text{skip} && \} \text{ termination} \\ &| \mathbb{X} && \} \text{ process variable} \\ &| \text{fix } \mathbb{X}. P && \} \text{ recursion} \\ &| P; P && \} \text{ seq. composition} \end{aligned}$$

Let us comment on the most peculiar constructs. With $\text{listen}(L)$ an app listens on a set of services L whose changes may trigger the app to execute. This is a blocking construct as it progresses only when at least one of the services listed in L changes. L is formally defined as follows:

$$\begin{aligned} \text{VarList} \ni L &::= L; L && \} \text{ services list} \\ &| x && \} \text{ service} \end{aligned}$$

The construct $x \leftarrow e$ sets the local variable x (the local view of the global service x) with the value obtained by the evaluation of an expression e . Note that, in the expression e we may have both readings on local variables y , simply denoted with y , and readings of global variables y , denoted with $\text{read}(y)$. Thus, in the assignment $x \leftarrow \text{read}(x) + y$ the local copy of the service variable x is updated with the summation between the up-to-date value of the global service x (taken from the cloud) and the value read from the local copy of the service y . The construct $\text{update}(x)$ updates the value of the service x in the global store with its current value in the local store. The process $\text{fix } \mathbb{X}. P$ is the standard construct to denote recursion.

An app is a process silently running in background until a trigger occurs. This latter fires the app payload, consisting of a sequence of actions (potentially dispatched after the execution of some code). Technically speaking, the process running in an app is a recursive process of the form: $\text{fix } \mathbb{X}. \text{listen}(L); \text{payload}$. Intuitively, our apps keep listening on a number of cloud services: when at least one these services changes, the app executes its payload. The payload consists in performing a number of activities, such as checking the state of some cloud service x via the $\text{read}(x)$ expression, and updating one or more cloud services via the $\text{update}(x)$ construct.

Actually, the syntax proposed for the code of our apps is a bit too permissive with respect to our intentions. We could rule out ill-formed apps with a simple type systems. However, for the sake of simplicity, we prefer to provide the following definition.

Definition 1 (Well-Formedness). An app $\text{id}[D \bowtie P]$ is *well-formed* if the following conditions are satisfied:

- P is of the form $\text{fix } \mathbb{X}. \text{listen}(L); Q$;
- x appears in $\text{listen}(L)$ only if x^R occurs in D ;
- the payload Q does not contain listeners;
- $\text{read}(x)$ appears in Q only if x^R occurs in D ;
- $\text{update}(x)$ appears in Q only if x^W occurs in D .

A system is *well-formed* only if its apps are well-formed.

Hereafter, we will always work with well-formed systems.

Let us provide two simple examples to describe how we can model IoT apps in CaITApp .

Example 1. Consider the following two apps. Tw2Fb reposts on Facebook messages received on Twitter. Similarly, when there is a new post on Facebook, the app Fb2Ld publishes the post on LinkedIn. In this case, we have three logical services: tw , for Twitter, fb , for Facebook and ld , for LinkedIn. The apps are formalized in our language as follows:

$$\text{Tw2Fb}[\text{tw}^R; \text{fb}^W \bowtie \text{fix } \mathbb{X}. \text{listen}(\text{tw}); \text{pld1}]$$

where $\text{pld1} \stackrel{\text{def}}{=} \text{tw} \leftarrow \text{read}(\text{tw}); \text{fb} \leftarrow \text{tw}; \text{update}(\text{fb}); \mathbb{X}$, and

$$\text{Fb2Ld}[\text{fb}^R; \text{ld}^W \bowtie \text{fix } \mathbb{X}. \text{listen}(\text{fb}); \text{pld2}]$$

where $\text{pld2} \stackrel{\text{def}}{=} \text{ld} \leftarrow \text{read}(\text{fb}); \text{ld} \leftarrow \text{fb}; \text{update}(\text{ld}); \mathbb{X}$.

Example 2. Consider the following two apps. When smoke is detected, the app SmokeAlarm should fire the smoke alarm and turn on the lights. If a given heat threshold is reached, then the app Sprinks will open the water valve to activate fire sprinkles. For that we assume five logical services: smoke , reporting the presence of smoke, heat , reporting the heat level, waterV , controlling the water valve, alarm , controlling the smoke alarm, and lights , managing the lights. The apps are formalized in our language as

$$\text{SmokeAlarm}[\text{smoke}^R; \text{alarm}^W; \text{lights}^W \bowtie \text{fix } \mathbb{X}. \text{listen}(\text{smoke}); P3]$$

where:

$$\begin{aligned} P3 &\stackrel{\text{def}}{=} \text{smoke} \leftarrow \text{read}(\text{smoke}); \\ &\text{if } (\text{smoke} = \text{yes}) \text{ then } \{ \\ &\quad \text{alarm} \leftarrow \text{On}; \\ &\quad \text{lights} \leftarrow \text{On}; \\ &\quad \text{update}(\text{alarm}, \text{lights}) \\ &\}; \mathbb{X} \end{aligned}$$

and $\text{Sprinks}[\text{heat}^R; \text{waterV}^W \bowtie \text{fix } \mathbb{X}. \text{listen}(\text{heat}); P4]$, where:

$$\begin{aligned} P4 &\stackrel{\text{def}}{=} \text{heat} \leftarrow \text{read}(\text{heat}); \\ &\text{if } (\text{heat} \geq 45) \text{ then } \{ \\ &\quad \text{waterV} \leftarrow \text{Open}; \text{update}(\text{waterValve}) \\ &\}; \mathbb{X} \end{aligned}$$

$$\begin{array}{c}
\text{(StopListening)} \frac{L = x_1; \dots; x_n \quad \exists i \in [1, n]. \mathfrak{G}(x_i) \neq \phi(x_i)}{\langle \mathfrak{G}, \phi \rangle \triangleright \text{listen}(L) \xrightarrow{\tau} \langle \mathfrak{G}, \phi \rangle \triangleright \text{skip}} \\
\text{(SetLocal)} \frac{\llbracket e \rrbracket(\mathfrak{G}, \phi) = v}{\langle \mathfrak{G}, \phi \rangle \triangleright x \leftarrow e \xrightarrow{\tau} \langle \mathfrak{G}, \phi[x \leftarrow v] \rangle \triangleright \text{skip}} \\
\text{(SkipUpdate)} \frac{\mathfrak{G}(x) = \phi(x)}{\langle \mathfrak{G}, \phi \rangle \triangleright \text{update}(x) \xrightarrow{\tau} \langle \mathfrak{G}, \phi \rangle \triangleright \text{skip}} \\
\text{(Update)} \frac{\mathfrak{G}(x) \neq \phi(x) \quad \phi(x) = v}{\langle \mathfrak{G}, \phi \rangle \triangleright \text{update}(x) \xrightarrow{x!v} \langle \mathfrak{G}[x \leftarrow v], \phi \rangle \triangleright \text{skip}} \\
\text{(IfTrue)} \frac{\llbracket b \rrbracket(\mathfrak{G}, \phi) = \mathbf{tt}}{\langle \mathfrak{G}, \phi \rangle \triangleright \text{if } b \text{ then } \{P_1\} \text{ else } \{P_2\} \xrightarrow{\tau} \langle \mathfrak{G}, \phi \rangle \triangleright P_1} \\
\text{(IfFalse)} \frac{\llbracket b \rrbracket(\mathfrak{G}, \phi) = \mathbf{ff}}{\langle \mathfrak{G}, \phi \rangle \triangleright \text{if } b \text{ then } \{P_1\} \text{ else } \{P_2\} \xrightarrow{\tau} \langle \mathfrak{G}, \phi \rangle \triangleright P_2} \\
\text{(Fix)} \frac{-}{\langle \mathfrak{G}, \phi \rangle \triangleright \text{fix } X \bullet P \xrightarrow{\tau} \langle \mathfrak{G}, \phi \rangle \triangleright P \{ \text{fix } X \bullet P / X \}} \\
\text{(SeqSkip)} \frac{-}{\langle \mathfrak{G}, \phi \rangle \triangleright \text{skip}; P \xrightarrow{\tau} \langle \mathfrak{G}, \phi \rangle \triangleright P} \\
\text{(Seq)} \frac{\langle \mathfrak{G}, \phi \rangle \triangleright P_1 \xrightarrow{\lambda} \langle \mathfrak{G}', \phi' \rangle \triangleright P'_1}{\langle \mathfrak{G}, \phi \rangle \triangleright P_1; P_2 \xrightarrow{\lambda} \langle \mathfrak{G}', \phi' \rangle \triangleright P'_1; P_2}
\end{array}$$

TABLE I
LABELED TRANSITION SEMANTICS FOR PROCESSES

We end this section with a couple of *notations and abbreviations* that will be used in the rest of the paper. We will write $\text{update}(x_1, x_2, \dots, x_n)$ as an abbreviation for the sequential update of the global variables x_1, x_2, \dots, x_n , namely: $\text{update}(x_1); \text{update}(x_2); \dots; \text{update}(x_n)$. We write $\text{if } b \text{ then } \{P\}$ as an abbreviation for $\text{if } b \text{ then } \{P\} \text{ else } \{\text{skip}\}$.

A. Labeled Transition Semantics

IoT Apps are simple applications interacting with physical and logical services that can be accessed only via a *cloud platform*, that we call *global store*, denoted with $\mathfrak{G} \in \mathbb{S}$, where $\mathbb{S} \stackrel{\text{def}}{=} \text{Service} \rightarrow \text{Value}$ is the set of all partial functions from services to values. Every app $\text{id}[D \rtimes P]$ retains a *local view* of the cloud platform that must be consistent with the app's declaration D , meaning that the domain of the *local store* of id consists of all and only those services declared in D . Changes in the global store are shared with all parallel apps of the system associated to the same user/account; however, these modifications do not directly affect the apps' local view of the store. Indeed, a local store can be modified only explicitly by its related app payload.

Since our syntax distinguishes between processes and systems of apps, in our labeled transition semantics we have two different kinds of transitions: one for processes and a second one for systems.

$$\begin{array}{c}
\text{(App)} \frac{\mathfrak{L}(\text{id}) = \phi \quad \langle \mathfrak{G}, \phi \rangle \triangleright P \xrightarrow{\tau} \langle \mathfrak{G}, \phi' \rangle \triangleright P'}{\langle \mathfrak{G}, \mathfrak{L} \rangle \triangleright \text{id}[D \rtimes P] \xrightarrow{\tau} \langle \mathfrak{G}, \mathfrak{L}[\text{id} \leftarrow \phi'] \rangle \triangleright \text{id}[D \rtimes P']} \\
\text{(AppUpdate)} \frac{\mathfrak{L}(\text{id}) = \phi \quad \langle \mathfrak{G}, \phi \rangle \triangleright P \xrightarrow{x!v} \langle \mathfrak{G}', \phi \rangle \triangleright P'}{\langle \mathfrak{G}, \mathfrak{L} \rangle \triangleright \text{id}[D \rtimes P] \xrightarrow{\text{id}:x!v} \langle \mathfrak{G}', \mathfrak{L} \rangle \triangleright \text{id}[D \rtimes P']} \\
\text{(EnvChange)} \frac{-}{\langle \mathfrak{G}, \mathfrak{L} \rangle \triangleright S \xrightarrow{x?v} \langle \mathfrak{G}[x \leftarrow v], \mathfrak{L} \rangle \triangleright S} \\
\text{(ParLeft)} \frac{\langle \mathfrak{G}, \mathfrak{L} \rangle \triangleright S_1 \xrightarrow{\alpha} \langle \mathfrak{G}', \mathfrak{L}' \rangle \triangleright S'_1 \quad \alpha \in \{\tau, \text{id}:x!v\}}{\langle \mathfrak{G}, \mathfrak{L} \rangle \triangleright S_1 \parallel S_2 \xrightarrow{\alpha} \langle \mathfrak{G}', \mathfrak{L}' \rangle \triangleright S'_1 \parallel S_2} \\
\text{(ParRight)} \frac{\langle \mathfrak{G}, \mathfrak{L} \rangle \triangleright S_2 \xrightarrow{\alpha} \langle \mathfrak{G}', \mathfrak{L}' \rangle \triangleright S'_2 \quad \alpha \in \{\tau, \text{id}:x!v\}}{\langle \mathfrak{G}, \mathfrak{L} \rangle \triangleright S_1 \parallel S_2 \xrightarrow{\alpha} \langle \mathfrak{G}', \mathfrak{L}' \rangle \triangleright S_1 \parallel S'_2}
\end{array}$$

TABLE II
LABELED TRANSITION SEMANTICS FOR SYSTEMS

In Table I we provide the transition rules for *process configurations* of the form

$$\langle \mathfrak{G}, \phi \rangle \triangleright P \xrightarrow{\lambda} \langle \mathfrak{G}', \phi' \rangle \triangleright P',$$

where $\mathfrak{G} \in \mathbb{S}$ denotes the global store while $\phi \in \mathbb{S}$ is the local store in which the process P is running. The labels $\lambda \in \{\tau, x!v\}$ denote: *nonobservable* actions and *observable* modifications (writings) of a global service x , respectively. We assume a standard *evaluation semantics* for expressions $\llbracket e \rrbracket \in \mathbb{S} \times \mathbb{S} \rightarrow \text{Value}$, inductively defined on the structure of the expression e . We omit the details of its definition. Here, it is only important to notice that we have one expression for looking-up variables in the local store, $\llbracket x \rrbracket(\mathfrak{G}, \phi) \stackrel{\text{def}}{=} \phi(x)$, and a different expression for looking-up variables in the global store: $\llbracket \text{read}(x) \rrbracket(\mathfrak{G}, \phi) \stackrel{\text{def}}{=} \mathfrak{G}(x)$.

Now, let us comment on the most interesting rules of Table I. The construct $\text{listen}(L)$ is a blocking operator waiting for changes in the cloud on (at least on of) those services contained in L . The semantics of the $\text{listen}(L)$ operator is formalized by means of the rule (StopListening) . The transition rule (SetLocal) serves to update the local store via an assignment to (the local copy of the) service x ; this assignment will affect the global service x only if followed by an $\text{update}(x)$. By an application of the rule (Update) the construct $\text{update}(x)$ modifies the value associated to the service x in the global store with the value recorded in the local store. The remaining rules are standard.

In Table II we provide the transition rules for *system configurations* of the form

$$\langle \mathfrak{G}, \mathfrak{L} \rangle \triangleright S \xrightarrow{\alpha} \langle \mathfrak{G}', \mathfrak{L}' \rangle \triangleright S',$$

where $\mathfrak{G} \in \mathbb{S}$ denotes the global store, whereas $\mathfrak{L} \in \mathcal{I} \rightarrow \mathbb{S}$ is a map associating to any app identifier its local store. The labels $\alpha \in \{\tau, \text{id}:x!v, x?v\}$ denote: *nonobservable* actions, *observable* modifications (writings) made by the applet id on a global service x , and *observable* changes on a global

service x made by the external environment, respectively. In the following, we will write $\mathcal{C} \xrightarrow{\alpha} \mathcal{C}'$ to denote a transition between system configurations belonging to the set Sconf of all possible system configurations.

Now, let us comment on the transition rules of Table II. Both rules (App) and (AppUpdate) serve to lift actions from processes to apps (and hence systems of apps). Observe that in case of updates on the cloud services, we are interested in annotating the label of the action with the name of the app performing the write operation; this will be useful when defining *safe cross-app interactions* (Definition 3). The rule (EnvChange) models changes in the cloud made by the external environment and affecting all apps. Thus, this action is not triggered by some app of the system and it can be fired nondeterministically at any moment. Rules (ParLeft) and (ParRight) are standard. Note that, for convenience, action $x?v$ is allowed only to complete systems as it does not propagate through parallel composition (it is not admitted in rules (ParLeft) and (ParRight)).

III. DEFINING SAFE CROSS-APP INTERACTIONS

In order to capture harmful interactions in systems of apps, we formalize a notion of *safe cross-app interaction* based on a bisimulation-based behavioral semantics for our systems.

Intuitively, two apps may interact with each other by acting on some common services in such a way that the state reached by those services is somehow inconsistent (think of a thermostat or a valve when activated by different apps designed with different specifications in mind). However, this is not the only way to yield unsafe interactions between two apps: an app A might interact with a second app B if the execution of some actions of A may affect services in the cloud whose changes may subsequently trigger activities of B . Those activities of B would not have been fired if A would not have modified the state of its services on the cloud.

A. Semantic characterization of safe cross-app interactions

In this section, we provide a semantic characterization of safe cross-app interaction based on some appropriate notion of bisimulation.

Intuitively, we would like to say that a system of apps S does not interact with a system R if the runtime behavior of R when running in parallel with S does not differ from its behavior when running in isolation. A bit more formally, along the lines of Focardi and Martinelli's *Generalized Non Deducibility on Composition (GNDC)* [21], we would like to say that a system S does not interact with a system R if

$$S \parallel R \simeq_S R$$

for some appropriate bisimilarity \simeq_S that hides those (observable) actions of S that modify the services in the cloud (the global store). Notice that the bisimilarity \simeq_S should only suppress the capability of the observer to detect writing actions on the cloud services executed by S ; however, these writings must be executed, so that indirect interactions via the cloud between the two systems can be observed if they trigger a nongenuine behavior in R .

Basically, in the scenario above, if bisimilarity breaks then the system S does interact with the correct execution of R in at least one of the following ways:

- The compound system $S \parallel R$ might have nongenuine traces containing observables (originating from the R component) that cannot be reproduced by R in isolation; here the interaction affects the *integrity* of the behavior of R .
- The system R might have execution traces containing observables that cannot be reproduced by the compound system $S \parallel R$ because they are prevented S ; this is a violation of the *availability* of the system R .

In order to formalize the concepts described above, we define a slight generalization of the *weak asynchronous bisimulation* [3] introduced for the asynchronous fragment of the π -calculus [22], [28]. In that bisimulation, input actions are made not observable because in an asynchronous setting the observer cannot directly observe them. Here, we intend to hide modifications on cloud enabled by the interacting system.

Consider a set H of *hidden actions*, $H \subseteq \mathcal{A} \setminus \{\tau\}$. Then, the following bisimulation compares two system configurations by observing all possible actions except those occurring in H .

Definition 2 (Hiding Bisimulation). Given a set of actions $H \subseteq \mathcal{A} \setminus \{\tau\}$, the symmetric relation $\mathcal{R} \subseteq \text{Sconf} \times \text{Sconf}$ is a *hiding bisimulation* parametric on H if and only if, whenever $\mathcal{C}_1 \mathcal{R} \mathcal{C}_2$ and $\mathcal{C}_1 \xrightarrow{\alpha} \mathcal{C}'_1$ we have the following:

- if $\alpha \notin H$ then $\mathcal{C}_2 \xrightarrow{\hat{\alpha}} \mathcal{C}'_2$, for some \mathcal{C}'_2 such that $\mathcal{C}'_1 \mathcal{R} \mathcal{C}'_2$;
- if $\alpha \in H$ then
 - either $\mathcal{C}_2 \xrightarrow{\hat{\alpha}} \mathcal{C}'_2$, for some \mathcal{C}'_2 such that $\mathcal{C}'_1 \mathcal{R} \mathcal{C}'_2$
 - or $\mathcal{C}_2 \xrightarrow{\tau^*} \mathcal{C}'_2$, for some \mathcal{C}'_2 such that $\mathcal{C}'_1 \mathcal{R} \mathcal{C}'_2$.

We say that two system configurations \mathcal{C}_1 and \mathcal{C}_2 are *hiding bisimilar* w.r.t. the set of actions H , written $\mathcal{C}_1 \approx_H \mathcal{C}_2$, if $\mathcal{C}_1 \mathcal{R} \mathcal{C}_2$ for some hiding bisimulation \mathcal{R} parametric on H .

Obviously, for $H = \emptyset$ we get the standard bisimulation.

In the following, for the sake of simplicity, given two system configurations $\langle \mathcal{G}, \mathcal{L} \rangle \triangleright S$ and $\langle \mathcal{G}, \mathcal{L} \rangle \triangleright R$, we will write

$$\langle \mathcal{G}, \mathcal{L} \rangle \triangleright S \approx_H R$$

as an abbreviation for $\langle \mathcal{G}, \mathcal{L} \rangle \triangleright S \approx_H \langle \mathcal{G}, \mathcal{L} \rangle \triangleright R$.

Now, we can use our hiding bisimilarity to formalize a *semantic-based* notion of safe cross-app interaction. As said before, our intention is to hide only those actions that may cause an update on the cloud. Thus, given an arbitrary system $S \stackrel{\text{def}}{=} \prod_{i=1}^n \text{id}_i[D_i \bowtie P_i]$, we define the set of possible actions of S that may modify the state of the cloud services:

$$\text{upd}(S) \stackrel{\text{def}}{=} \bigcup_{1 \leq i \leq n} \{\text{id}_i : x!v \mid x^w \in D_i\}.$$

In the following, we let \mathcal{L}_\perp be the function $\lambda \text{id}. \lambda x. \perp$ used to define an initial local environments for all apps in which all services are not initialized.

Definition 3 (Safe cross-app interaction). Let S and R be two systems of apps in CaITApp . We say that S is *noninteracting* with R when

$$\forall \mathcal{G} \in \mathbb{S}. \langle \mathcal{G}, \mathcal{L}_\perp \rangle \triangleright S \parallel R \approx_{H_S} R,$$

for $H_S \stackrel{\text{def}}{=} \text{up}\delta(S)$. We say that the two systems S and R *do not interact with each other* if in addition to the previous requirement it holds that

$$\forall \mathcal{G} \in \mathbb{S}. \langle \mathcal{G}, \mathcal{L}_\perp \rangle \triangleright S \parallel R \approx_{H_R} S,$$

for $H_R \stackrel{\text{def}}{=} \text{up}\delta(R)$.

Example 3. Consider the two apps Tw2Fb and Fb2Ld introduced in Example 1. Here, we have a potentially unwanted interaction between the two apps as Tw2Fb may trigger Fb2Ld : a message of Twitter will be posted on Facebook by Tw2Fb , then the app Fb2Ld will post that message on LinkedIn. In fact, according to Definition 3, the app Tw2Fb may interact with the app Fb2Ld as:

$$\exists \mathcal{G} \in \mathbb{S}. \langle \mathcal{G}, \mathcal{L}_\perp \rangle \triangleright \text{Tw2Fb} \parallel \text{Fb2Ld} \not\approx_H \text{Fb2Ld}$$

for $H = \{\text{Tw2Fb:tw!}v \mid v \in \text{Value}\}$. It is easy to see that on the left-hand-side compound system the app Fb2Ld might do a writing on the cloud service LinkedIn that would not occur if the app was running in isolation.

Example 4. Consider the two apps SmokeAlarm and Sprinks introduced in Example 2. According to Definition 3, the two apps do not interact with each other, as for any $\mathcal{G} \in \mathbb{S}$ we have the following:

- $\langle \mathcal{G}, \mathcal{L}_\perp \rangle \triangleright \text{SmokeAlarm} \parallel \text{Sprinks} \approx_{H_1} \text{SmokeAlarm}$, with $H_1 \stackrel{\text{def}}{=} \{\text{Sprink:waterValve!}v \mid v \in \text{Value}\}$;
- $\langle \mathcal{G}, \mathcal{L}_\perp \rangle \triangleright \text{SmokeAlarm} \parallel \text{Sprinks} \approx_{H_2} \text{Sprinks}$, with $H_2 \stackrel{\text{def}}{=} \{\text{SmokeAlarm:alarm!}v, \text{SmokeAlarm:lights!}v\}$, for any $v \in \text{Value}$.

In the example above the two apps do not interact with each other simply because they work on different services. However, according to Definition 3, two apps may be non-interacting even if they write on the same services, provided that no causalities exist among the two writings.

Example 5. Suppose to have an app SimPres that simulates the presence of the user when it is not at home during the night, turning on lights for 10 minutes every half an hour. Then, consider a second app eSaver turning off lights during the day to save energy whenever there is no motion for at least 5 minutes. The two apps are defined as follows:

$$\text{SimPres}[\text{user}^R; \text{time}^R; \text{lights}^W \bowtie \text{fix } \mathbb{X} \bullet \text{listen}(\text{user}; \text{time}); P_5],$$

where:

$$P_5 \stackrel{\text{def}}{=} \text{user} \leftarrow \text{read}(\text{user}); \\ \text{if } (0 < \text{read}(\text{time.H}) < 7 \wedge \text{user} = \text{away}) \text{ then } \{ \\ \quad \text{if } \text{read}(\text{time.M}) = 30 \text{ then } \{ \\ \quad \quad \text{lights} \leftarrow \text{On10minsOff}; \text{update}(\text{lights}) \\ \quad \quad \} \\ \quad \} \\ \}; \mathbb{X}$$

$\text{eSaver}[\text{none}^R; \text{time}^R; \text{lights}^{\text{RW}} \bowtie \text{fix } \mathbb{X} \bullet \text{listen}(\text{none}; \text{lights}); P_6]$, where:

$$P_6 \stackrel{\text{def}}{=} \text{none} \leftarrow \text{read}(\text{none}); \\ \text{if } (8 < \text{read}(\text{time.H}) < 18) \text{ then } \{ \\ \quad \text{if } (\text{none} \geq 5\text{mins} \wedge \text{lights} = \text{On}) \text{ then } \{ \\ \quad \quad \text{lights} \leftarrow \text{Off}; \text{update}(\text{lights}) \\ \quad \quad \} \\ \quad \} \\ \}; \mathbb{X}$$

Notice that that there is no interaction between these two apps, even if they write on the same global service lights . Actually, those writings occur in different periods of the day and can never interact. Thus, according to Definition 3, for any $\mathcal{G} \in \mathbb{S}$ we have the following:

- $\langle \mathcal{G}, \mathcal{L}_\perp \rangle \triangleright \text{SimPres} \parallel \text{eSaver} \approx_{H_1} \text{eSaver}$, with $H_1 \stackrel{\text{def}}{=} \{\text{SimPres:lights!}v \mid v \in \text{Value}\}$;
- $\langle \mathcal{G}, \mathcal{L}_\perp \rangle \triangleright \text{SimPres} \parallel \text{eSaver} \approx_{H_2} \text{SimPres}$, with $H_2 \stackrel{\text{def}}{=} \{\text{eSaver:lights!}v \mid v \in \text{Value}\}$.

B. A proof technique for safe cross-app interaction

Although the notion of safe cross-app interaction in Definition 3 is very intuitive, it is actually quite hard to verify due to the universal quantification over all possible global stores.

In this section, we provide *syntactic conditions*, easy to verify, that allow us to enforce the semantic condition of safe cross-app interaction. In order to do that, we have to formally specify: (i) what are the potential actions that an app may perform, (ii) what are the services whose changes may trigger an app.

In our calculus CaITApp , the actions potentially performed by an app $\text{id}[D \bowtie P]$ are given by the services declared in write modality.

Definition 4 (Actions). Given an app $\text{id}[D \bowtie P]$, we define $\text{act}(\text{id}) \stackrel{\text{def}}{=} \{x \in \text{Service} \mid x^W \in D\}$. More generally, in a system of apps $S \stackrel{\text{def}}{=} \prod_{i=1}^n \text{id}_i[D_i \bowtie P_i]$ we define $\text{act}(S) \stackrel{\text{def}}{=} \bigcup_{1 \leq i \leq n} \text{act}(\text{id}_i)$.

Similarly, the triggers of an app $\text{id}[D \bowtie P]$ are given by the services on which the app currently listen or make a read from the global store, namely the services declared in read modality.

Definition 5 (Triggers). Given an app $\text{id}[D \bowtie P]$, we define $\text{trg}(\text{id}) \stackrel{\text{def}}{=} \{x \in \text{Service} \mid x^R \in D\}$. More generally, in a system of apps $S \stackrel{\text{def}}{=} \prod_{i=1}^n \text{id}_i[D_i \bowtie P_i]$ we define $\text{trg}(S) \stackrel{\text{def}}{=} \bigcup_{1 \leq i \leq n} \text{trg}(\text{id}_i)$.

Now, everything is in place to provide a syntactic condition for safe cross-app interaction, where a system S is said not to interact with a system R when:

- the two systems do not write on common cloud services;
- the execution of S may not trigger any app of R .

Formally,

Definition 6 (Syntax-based safe cross-app interaction). The system S is said to be *syntactically noninteracting* with the

system R , written $S \rightarrow R$, when both the following conditions hold:

- 1) $\text{act}(S) \cap \text{act}(R) = \emptyset$;
- 2) $\text{act}(S) \cap \text{trg}(R) = \emptyset$.

More generally, we say that the two systems S and R are *syntactically noninterfering w.r.t. each other*, written $S \leftrightarrow R$, when besides the two conditions above it also holds:

- 3) $\text{trg}(S) \cap \text{act}(R) = \emptyset$.

Now, if we consider the apps in Examples 1, 2 and 5, it is easy to verify that:

- $\text{Fb2Ld} \rightarrow \text{Tw2Fb}$ holds;
- $\text{Tw2Fb} \rightarrow \text{Fb2Ld}$ does not hold because $\text{act}(\text{Tw2Fb}) \cap \text{trg}(\text{Fb2Ld}) \neq \emptyset$;
- $\text{SmokeAlarm} \leftrightarrow \text{Sprinks}$ holds;
- $\text{SimPres} \leftrightarrow \text{eSaver}$ does not hold because they can both modify the service `lights`, that is $\text{act}(\text{SimPres}) \cap \text{act}(\text{eSaver}) = \{\text{lights}\} \neq \emptyset$.

Thus, Definition 6 provides an easy-to-verify syntactic sufficient condition for semantic-based condition.

Theorem 1 (Soundness). *Let S and R be two systems of apps in CaITApp . Let $H_S \stackrel{\text{def}}{=} \text{up}\delta(S)$ and $H_R \stackrel{\text{def}}{=} \text{up}\delta(R)$. Then:*

- $S \rightarrow R$ implies $\forall \mathcal{G} \in \mathbb{S}. \langle \mathcal{G}, \mathcal{L}_\perp \rangle \triangleright S \parallel R \approx_{H_S} R$;
- $S \leftrightarrow R$ implies
 - $\forall \mathcal{G} \in \mathbb{S}. \langle \mathcal{G}, \mathcal{L}_\perp \rangle \triangleright S \parallel R \approx_{H_S} R$
 - $\forall \mathcal{G} \in \mathbb{S}. \langle \mathcal{G}, \mathcal{L}_\perp \rangle \triangleright S \parallel R \approx_{H_R} S$.

The details of the proof can be found in the Appendix.

Obviously, Definition 6 provides us with a sufficient but not necessary condition to derive soundness for cross-app interactions, as shown, for instance, by the two apps `SimPres` and `eSaver` in Example 5.

IV. IMPLICIT INTERACTIONS AND PRIORITIES

In this section, we study: (i) the challenge posed by *implicit cross-app interactions*, and (ii) the possibility of having *priorities* between different services.

The former arises whenever two services, *e.g.*, temperature and thermostat, are semantically related, though they differ syntactically. This may lead to both the semantic condition and the enforcement mechanism deeming an interaction as safe, while this is not the case in practice. We propose an extension of our language semantics, as well as a semantic condition and a syntactic one to reason about such cases.

The latter may be useful when our semantic condition in Definition 3 becomes too restrictive for use cases where the end user is willing to accept some interactions on specific services, *e.g.*, social networks like Facebook, while avoiding interactions on other services, *e.g.*, the temperature. By leveraging a lattice order of priorities between services, we discuss how our condition can be extended to enable such flexibility.

A. Countering implicit interactions

The semantic-based condition given in Definition 3 works quite well when dealing with logical services like Facebook or Twitter as in Example 1. However, when stepping to physical

services, *i.e.*, services affecting the physical environment, such as the temperature of a room, we may end up accepting as safe a system of apps in the presence of some kind *implicit interactions*. Consider the example below.

Example 6. Let `Win` be an app managing the window of a room, depending on the temperature detected: when the temperature is above 22 degrees then the window must be opened. Formally,

$$\text{Win}[\text{temp}^R; \text{win}^W \times \text{fix } \mathbb{X}. \text{listen}(\text{temp}); \text{pld7}]$$

with

$$\begin{aligned} \text{pld7} &\stackrel{\text{def}}{=} \text{temp} \leftarrow \text{read}(\text{temp}); \\ &\text{if } (\text{temp} > 22) \text{ then } \{\text{win} \leftarrow \text{Open}; \text{update}(\text{win})\}; \\ &\mathbb{X} \end{aligned}$$

Now, suppose to have a second app `Therm`, managing the thermostat of the room, such that when the temperature is below 17 degrees the thermostat is set to rise the temperature by 3 degrees. Formally,

$$\text{Therm}[\text{temp}^R; \text{therm}^W \times \text{fix } \mathbb{X}. \text{listen}(\text{temp}); \text{pld8}]$$

with

$$\begin{aligned} \text{pld8} &\stackrel{\text{def}}{=} \text{temp} \leftarrow \text{read}(\text{temp}); \\ &\text{if } (\text{temp} < 17) \text{ then } \{\text{therm} \leftarrow +3; \text{update}(\text{therm})\}; \\ &\mathbb{X} \end{aligned}$$

When running these two apps in parallel, we may have an *implicit interaction*, as the app `Therm` may indirectly trigger the app `Win`. This is because, we know, out of band, that the temperature of the room should somehow change according to the changes made on the thermostat of the room.

However, since this out-of-band information is not considered by our formalization, according to Definition 3 we would have a kind of *false negative* as the app `Therm` is not directly interacting with the app `Win`. Formally, for $H \stackrel{\text{def}}{=} \text{up}\delta(\text{Therm}) = \{\text{Therm}:\text{therm}!v \mid v \in \text{Value}\}$, we have:

$$\langle \mathcal{G}, \mathcal{L}_\perp \rangle \triangleright \text{Win} \parallel \text{Therm} \approx_H \text{Win},$$

for any $\mathcal{G} \in \mathbb{S}$.

Note that *causality dependencies* between services, such as those asserting that thermostat changes may affect the temperature, are not part of the specification of an app (or of a system of apps). This information comes from the physics of the real system managed via apps. Thus, by no means we can capture this kind of *implicit interactions* unless we have information about causality dependencies.

However, we can assume that, when designing our system of apps we actually get, out of band, a set of causality dependencies to improve the precision of our analysis ruling out a number of false negatives. For the sake of simplicity, we define a *dependency policy* as a binary relation $\Delta \subseteq \text{Service} \times \text{Service}$ such that $(x, y) \in \Delta$ when the service y may be affected by changes occurring at the service x . Clearly, dependencies can be composed, hence we will consider the reflexive and

transitive closure of Δ in order to capture all possible dependencies associated to a service. We write $\text{clo}(\Delta, x)$ to denote the reflexive and transitive closure of the dependency policy Δ w.r.t. the service x . More generally, given a set of services $X \subseteq \text{Service}$ we define $\text{clo}(\Delta, X) \stackrel{\text{def}}{=} \bigcup_{x \in X} \text{clo}(\Delta, x)$.

Here, it is important to notice that when we act on the thermostat of the room we actually do not know exactly how the temperature will change (again, this depends on the physics, *e.g.*, on the wall isolation of the heated room). Thus, the dependency policy Δ records only abstract information relating pairs of services. More precisely, if $(x, y) \in \Delta$ we may assume that each time the service x changes on the cloud then the service y can be somehow affected. We represent this abstract information by means of *nondeterministic updates* assigning to y the special value $*$, meaning “any value”. Ideally, the special value $*$ satisfies any boolean expression containing it. For instance, $* \leq n$ is true for any n .

Now, using this extra out-of-band information Δ on the causality dependency between services, we can easily define a labeled transitions semantics $\xrightarrow{\alpha}_{\Delta}$, parametric on the set Δ :

- $\mathcal{C}_1 \xrightarrow{\alpha}_{\Delta} \mathcal{C}_2$ if $\mathcal{C}_1 \xrightarrow{\alpha} \mathcal{C}_2$ is derived by an application of any rule of Table I different from Update ;
- rule Update is replaced by the following transition rule:

$$\frac{\mathcal{G}(x) \neq \phi(x) \quad \phi(x) = v \quad \text{clo}(\Delta, x) = \{y_1, \dots, y_n\}}{\langle \mathcal{G}, \phi \rangle \triangleright \text{update}(x) \xrightarrow{x!v}_{\Delta} \langle \mathcal{G}[x \leftarrow v, y_1 \leftarrow *, \dots, y_n \leftarrow *], \phi \rangle \triangleright \text{skip}}$$

Now, we can refine Definition 3 making it parametric on a dependency policy Δ . Basically, we use our hiding bisimilarity defined on top of the parametric LTS $\xrightarrow{\alpha}_{\Delta}$, denoted with $\hat{\approx}_H$. In this manner, we can rely on the dependency policy Δ to capture false negatives due to implicit interactions: any change on a service x affects any service in the set $\text{clo}(\Delta, x)$ via nondeterministic assignments that will always trigger apps listening at these services.

Definition 7 (Safe cross-app interaction under dependencies). Let Δ be a dependency policy. Let S and R be two systems of apps in CaITApp . We say that S is *noninteracting with R under Δ* when

$$\forall \mathcal{G} \in \mathbb{S}. \langle \mathcal{G}, \mathcal{L}_{\perp} \rangle \triangleright S \parallel R \hat{\approx}_{H_S} R$$

where $H_S \stackrel{\text{def}}{=} \text{upd}(S)$. We say that the two systems S and R *do not interact with each other under Δ* if in addition to the requirement above we have

$$\forall \mathcal{G} \in \mathbb{S}. \langle \mathcal{G}, \mathcal{L}_{\perp} \rangle \triangleright S \parallel R \hat{\approx}_{H_R} S$$

where $H_R \stackrel{\text{def}}{=} \text{upd}(R)$.

Now, in order to provide a consistent reformulation of Theorem 1 to capture semantics-based noninteraction parametric on a dependency policy Δ , we have to reformulate Definition 6 to take into account the presence of the dependency policy.

Definition 8 (Syntax-based safe cross-app interaction under dependencies). Let Δ be a dependency policy. The system S is said to be *syntactically noninteracting under Δ* with the

system R , written $S \hat{\approx}_{\Delta} R$, when both the following conditions hold:

- 1) $\text{act}(S) \cap \text{act}(R) = \emptyset$;
- 2) $\text{clo}(\Delta, \text{act}(S)) \cap \text{trg}(R) = \emptyset$.

More generally, we say that the two systems S and R are *syntactically noninteracting w.r.t. each other under Δ* , written $S \hat{\approx}_{\Delta} R$, when besides the two conditions above we also have:

- 3) $\text{trg}(S) \cap \text{clo}(\Delta, \text{act}(R)) = \emptyset$.

Now, if we consider the apps in Example 6 with $\Delta = \{(\text{therm}, \text{temp})\}$ we have:

- $\text{Win} \hat{\approx}_{\Delta} \text{Therm}$ holds;
- $\text{Therm} \not\hat{\approx}_{\Delta} \text{Win}$ does not hold because $\text{trg}(\text{Win}) = \{\text{temp}\}$, $\text{clo}(\Delta, \text{act}(\text{Therm})) = \{\text{therm}, \text{temp}\}$ and hence $\text{clo}(\Delta, \text{act}(\text{Therm})) \cap \text{trg}(\text{Win}) \neq \emptyset$.

Finally, we can reformulate Theorem 1 to prove that Definition 8 provides a sufficient condition to capture semantic-based noninteraction under a given dependency policy Δ .

Theorem 2 (Soundness under dependencies). *Let Δ be a dependency policy. Let S and R be two systems of apps in CaITApp . Let $H_S \stackrel{\text{def}}{=} \text{upd}(S)$ and $H_R \stackrel{\text{def}}{=} \text{upd}(R)$. Then:*

- $S \hat{\approx}_{\Delta} R$ implies $\forall \mathcal{G} \in \mathbb{S}. \langle \mathcal{G}, \mathcal{L}_{\perp} \rangle \triangleright S \parallel R \hat{\approx}_{H_S} R$;
- $S \hat{\approx}_{\Delta} R$ implies
 - $\forall \mathcal{G} \in \mathbb{S}. \langle \mathcal{G}, \mathcal{L}_{\perp} \rangle \triangleright S \parallel R \hat{\approx}_{H_S} R$
 - $\forall \mathcal{G} \in \mathbb{S}. \langle \mathcal{G}, \mathcal{L}_{\perp} \rangle \triangleright S \parallel R \hat{\approx}_{H_R} S$.

The details of the proof can be found in the Appendix.

Thanks to Theorem 2, for $\Delta = \{(\text{therm}, \text{temp})\}$, we can now correctly capture the semantic-based interaction between the apps of Example 6 as there is $\mathcal{G} \in \mathbb{S}$ such that

$$\langle \mathcal{G}, \mathcal{L}_{\perp} \rangle \triangleright \text{Win} \parallel \text{Therm} \not\hat{\approx}_H \text{Win}$$

for $H \stackrel{\text{def}}{=} \text{upd}(\text{Therm}) = \{\text{Therm:therm!}v \mid v \in \text{Value}\}$.

B. Increasing flexibility via priorities

We envision use cases where users are aware of potential interactions on some services, while disallowing interactions on other services. For instance, an app managing the fire alarm may have higher priority than an app posting messages to Facebook. This relative importance of services can be specified with a lattice of service priorities. Suppose to have a *priority policy* Π , which associates a *priority level* p , taken from some complete lattice $\langle \text{PL}, \sqsubseteq \rangle$, with each service used by our system of apps.

Now, we can extend our condition of hiding bisimilarity to define safety of cross-app interference up to a given priority level. In fact, given a priority level p , we can formally ensure that two systems have the same behaviour only when looking at observables on services at priority level greater than, or equal to p .

Definition 9 (Safe cross-app interaction up to priorities). Let S and R be two systems of apps in CaITApp . We say that S is *noninteracting with R up to the priority level p* when:

$$\forall \mathcal{G} \in \mathbb{S}. \langle \mathcal{G}, \mathcal{L}_{\perp} \rangle \triangleright S \parallel R \approx_{H_p} R$$

where $H_p \stackrel{\text{def}}{=} \text{up}\partial(S) \cup \{\text{id}:x!v \in \text{up}\partial(R) \mid p \not\sqsubseteq \Pi(x)\}$.

Note that when Π maps each service to the same priority level, the definition above coincides with Definition 3. To enforce the condition above, we can easily define a syntax-based condition similar to the one given in Definition 6.

V. SECURING CROSS-APP INTERACTIONS

It is not hard to imagine that services accessed via an IoT platform may have different security clearance. For instance, a service to access a smart security camera should definitely not leak any kind of information to a second service that is used to share nice pictures among friends.

In this section, we assume a *security policy* Σ , which associates a *security level* σ , taken from some complete lattice $\langle SL, \preceq \rangle$, with each service used by our system of apps. For the sake of simplicity, in the examples, the security levels will simply be high (H), or *secret*, and low (L), or *public*, although the theory is developed for a generic complete lattice $\langle SL, \preceq \rangle$, of security levels. The goal is to achieve classical *noninterference* results stating that a system of apps is interference-free if its low-level services are not affected by changes occurring at its high-level services. Thus, information can safely flow from a service x to a service y whenever $\Sigma(x) \preceq \Sigma(y)$. In the following, we use \prec to denote the nonreflexive restriction of \preceq (i.e., $\delta \prec \sigma$ iff $\delta \preceq \sigma$ and $\delta \neq \sigma$).

As usual, a security policy induces an equivalence relation between stores. Given two (global) stores $\mathfrak{G}, \mathfrak{G}' \in \mathbb{S}$, we say that they are σ -equivalent, written $\mathfrak{G} \equiv_\sigma \mathfrak{G}'$, if they agree on the values associated to all services with security level lower than, or equal to, σ .

Definition 10 (σ -equivalent stores). Let $\mathfrak{G}, \mathfrak{G}' \in \mathbb{S}$ be two stores and $\sigma \in SL$ be a security level. We say that \mathfrak{G} and \mathfrak{G}' are σ -equivalent, written $\mathfrak{G} \equiv_\sigma \mathfrak{G}'$, whenever:

$$\forall x \in \text{Service} . \Sigma(x) \preceq \sigma \Rightarrow \mathfrak{G}(x) = \mathfrak{G}'(x).$$

Now, we can formalize a bisimulation-based notion of noninterference parametric on some security level $\sigma \in SL$. Intuitively, the runtime behavior at security level σ (or lower) of an interference-free system does not change when executed in two different σ -equivalent stores \mathfrak{G} and \mathfrak{G}' , though it may differ on services with security clearance higher than σ . Actually, in our notion of noninterference we consider σ -equivalent stores in $\mathbb{S}_\perp \stackrel{\text{def}}{=} \{\mathfrak{G} \in \mathbb{S} \mid \forall x \in \text{dom}(\mathfrak{G}) . \mathfrak{G}(x) \neq \perp\}$, as the mere initialization of an high-level service might activate a listener in an applet, thus leaking information about the *occurrence/presence* of a high event. We ignore presence leaks in order to increase premissiveness of our enforcement mechanism.

Our general notion of hiding bisimilarity can be used to hide (but not to suppress) actions involving changes affecting high-level services. In the following, with an abuse of notation, we extend Σ to assign security levels to system actions α ,

according to the cloud services affected by α . Thus, we define $\Sigma(\text{id}:x!v) = \Sigma(x)$ and $\Sigma(\alpha) = \perp$ for $\alpha \in \{\tau, x?v\}$.¹

However, in order to capture a semantic notion of non-interference that is not sensitive to information leaks due to *program termination*² we propose a modification of our hiding bisimilarity inspired by the *termination-insensitive i-bisimulation* proposed by Demange and Sands [16]. For this purpose, given a set of hidden actions H , we will write $\mathfrak{C}_1 \uparrow_H$ if and only if $\mathfrak{C}_1 \in D \stackrel{\text{def}}{=} \{\mathfrak{C} : (\exists \alpha \in \mathcal{A} . \mathfrak{C} \xrightarrow{\alpha} \mathfrak{C}') \wedge (\mathfrak{C} \xrightarrow{\alpha} \mathfrak{C}' \Rightarrow \alpha \in H \cup \{\tau\} \wedge \mathfrak{C}' \in D)\}$, that is \mathfrak{C}_1 belongs to the set of *high-level divergent configurations* that can always and only perform either τ -actions or high-level actions.

Definition 11 (Termination-Insensitive Hiding Bisimulation). Given a set of actions $H \subseteq \mathcal{A} \setminus \{\tau\}$, the symmetric relation $\mathcal{R} \subseteq \text{Sconf} \times \text{Sconf}$ is a *termination-insensitive hiding bisimulation* parametric on H if and only if, whenever $\mathfrak{C}_1 \mathcal{R} \mathfrak{C}_2$ and $\mathfrak{C}_1 \xrightarrow{\alpha} \mathfrak{C}'_1$ we have the following:

- if $\alpha \notin H$ then
 - either $\mathfrak{C}_2 \xrightarrow{\alpha} \mathfrak{C}'_2$, for some \mathfrak{C}'_2 such that $\mathfrak{C}'_1 \mathcal{R} \mathfrak{C}'_2$
 - or $\mathfrak{C}_2 \uparrow_H$
- if $\alpha \in H$ then
 - either $\mathfrak{C}_2 \xrightarrow{\alpha} \mathfrak{C}'_2$, for some \mathfrak{C}'_2 such that $\mathfrak{C}'_1 \mathcal{R} \mathfrak{C}'_2$
 - or $\mathfrak{C}_2 \xrightarrow{\tau^*} \mathfrak{C}'_2$, for some \mathfrak{C}'_2 such that $\mathfrak{C}'_1 \mathcal{R} \mathfrak{C}'_2$.

We say that two system configurations \mathfrak{C}_1 and \mathfrak{C}_2 are *termination-insensitive hiding bisimilar w.r.t.* the set of actions H , written $\mathfrak{C}_1 \approx_H^{\text{ti}} \mathfrak{C}_2$, if $\mathfrak{C}_1 \mathcal{R} \mathfrak{C}_2$ for some termination-insensitive hiding bisimulation \mathcal{R} parametric on H .

Definition 12 (σ -level noninterference). Let S be a system of apps and $H_\sigma \stackrel{\text{def}}{=} \{\alpha \in \mathcal{A} \mid \Sigma(\alpha) \not\preceq \sigma\}$ the set of actions with clearance greater than σ . We say that S is σ -level *interference-free* whenever:

$$\forall \mathfrak{G}, \mathfrak{G}' \in \mathbb{S}_\perp . \mathfrak{G} \equiv_\sigma \mathfrak{G}' \Rightarrow \langle \mathfrak{G}, \mathfrak{L}_\perp \rangle \triangleright S \approx_{H_\sigma}^{\text{ti}} \langle \mathfrak{G}', \mathfrak{L}_\perp \rangle \triangleright S.$$

Example 7. Consider the classic two-points lattice $\{L, H\}$, used for the system of apps $S \stackrel{\text{def}}{=} \text{Tw2Fb} \parallel \text{Fb2Ld}$ of Example 1 such that: $\Sigma(\text{tw}) = \Sigma(\text{fb}) = H$ and $\Sigma(\text{ld}) = L$. Obviously, the compound system is exposed to a security interference, as confidential information posted on tw will flow into the public service ld . In fact, it is not hard to find two L-equivalent global stores \mathfrak{G} and \mathfrak{G}' such that $\langle \mathfrak{G}, \mathfrak{L}_\perp \rangle \triangleright S \not\approx_{H_\perp}^{\text{ti}} \langle \mathfrak{G}', \mathfrak{L}_\perp \rangle \triangleright S$. This would not be the case if it was $\Sigma(\text{ld}) = H$. In that case, the bisimilarity would hold for any pair of L-equivalent stores.

Again, Definition 12 has a universal quantification on two global environments and then it requires the verification of a nontrivial bisimilarity. So, its verification is hard to achieve.

In order to provide a syntactic sufficient condition for security noninterference we resort to a type system, *parametric in the security policy* Σ , and inspired by the *flow-sensitive*

¹We recall that the action $x?v$ is not “originating” from our system but it denotes a modification of the service x made by the external environment.

²Termination leaks have well-known information-theoretic bounds [4], and they are usually ignored in order to increase permissiveness for static analysis that do not consider for program termination.

Typing rules for expressions

$$\begin{array}{c} (\text{Var}) \frac{\Gamma(\mathbf{x}) = \sigma}{\Gamma \vdash \mathbf{x} : \sigma} \quad (\text{Read}) \frac{\Sigma(\mathbf{x}) = \sigma}{\Gamma \vdash \text{read}(\mathbf{x}) : \sigma} \quad (\text{Expr}) \frac{\Gamma \vdash \mathbf{e}_1 : \sigma_1 \quad \Gamma \vdash \mathbf{e}_2 : \sigma_2}{\Gamma \vdash \mathbf{e}_1 \text{ op } \mathbf{e}_2 : \sigma_1 \sqcup \sigma_2} \end{array}$$

Typing rules for processes

$$\begin{array}{c} (\text{Skip}) \frac{-}{pc \vdash \Gamma\{\text{skip}\}\Gamma} \quad (\text{Assign}) \frac{\Gamma \vdash \mathbf{e} : \sigma}{pc \vdash \Gamma\{\mathbf{x} \leftarrow \mathbf{e}\}\Gamma[x \mapsto \sigma \sqcup pc]} \quad (\text{Fix}) \frac{pc \vdash \Gamma\{P\}\Gamma'}{pc \vdash \Gamma\{\text{fix } \mathbb{X} \bullet P\}\Gamma'} \\ (\text{Update}) \frac{\Gamma(\mathbf{x}) \preceq \Sigma(\mathbf{x})}{pc \vdash \Gamma\{\text{update}(\mathbf{x})\}\Gamma} \quad (\text{Pvar}) \frac{-}{pc \vdash \Gamma\{\mathbb{X}\}\Gamma} \quad (\text{Listen}) \frac{-}{pc \vdash \Gamma\{\text{listen}(L)\}\Gamma} \\ (\text{IfElse}) \frac{\Gamma \vdash \mathbf{b} : \sigma \quad pc \sqcup \sigma \vdash \Gamma_1\{P_1\}\Gamma_2 \quad pc \sqcup \sigma \vdash \Gamma_1\{P_2\}\Gamma_2}{pc \vdash \Gamma\{\text{if } \mathbf{b} \text{ then } \{P_1\} \text{ else } \{P_2\}\}\Gamma_2} \quad (\text{Seq}) \frac{pc \vdash \Gamma_1\{P_1\}\Gamma_2 \quad pc \vdash \Gamma_2\{P_2\}\Gamma_3}{pc \vdash \Gamma_1\{P_1; P_2\}\Gamma_3} \end{array}$$

Typing rules for systems and subtyping

$$\begin{array}{c} (\text{App}) \frac{pc \vdash \Gamma_1\{P\}\Gamma_2}{pc \vdash \text{id}[D \bowtie P]} \quad (\text{Par}) \frac{pc \vdash S_1 \quad pc \vdash S_2}{pc \vdash S_1 \parallel S_2} \quad (\text{Sub.Proc}) \frac{pc \vdash \Gamma_1\{P\}\Gamma_2 \quad pc' \preceq pc \quad \Gamma'_1 \preceq \Gamma_1 \quad \Gamma_2 \preceq \Gamma'_2}{pc' \vdash \Gamma'_1\{P\}\Gamma'_2} \end{array}$$

TABLE III
SECURITY TYPE SYSTEM

security type system of Hunt and Sands [23], adapted to our setting. In Table III we provide the typing rules.

Intuitively, a judgment of the form $pc \vdash S$ says that the system S does not contain information flows from services at security level higher than pc to services at security level lower than or equal to pc . Here, pc denotes the usual “program counter” level and serves to eliminate indirect information flows. We write $\vdash S$ to denote $pc \vdash S$ when pc is the least security level, *i.e.*, the bottom element of the lattice SL .

Since the syntax of our calculus is in two levels we also have a different kind of judgments for processes running inside an app: $pc \vdash \Gamma\{P\}\Gamma'$. Here, as in Hunt and Sands [23], Γ describes the security levels of services which hold before execution of P while Γ' describe the security level of those services after execution of P . Again, pc denotes the “program counter” level and the derivation rules ensure that only services which end up (in Γ') with types greater than or equal to pc may be changed by P .

Here, we wish to remark that, unlike batch-job noninterference models [37], a security information flow occurs in our setting only if a low-level service x is subject to an information flow and then x is “published” within the same app on the cloud via an update construct. In fact, the update operator is the only “exit gate” for potential (direct or indirect) information flows created within an app. This requires some care in the definition of the typing rule (Update). Basically, we impose that the update of a global service x is possible only if x is associated to an “original” security level (given by Σ) higher than or equal to the security level derived by its use (given by Γ) in the previous instructions within the app. On the other hand, we consider harmless those information flows that remain confined within an app because no update publishes their effects; this situations will not be ruled out by our type system. Finally, notice the difference between the two typing

rules (Var) and (Read) as they serve for typing accesses to local views of services and global services, respectively.

As expected, system (and process) well-typedness is preserved at runtime.

Proposition 1 (Subject reduction). *Let Σ be a security policy, S a system of apps, and $\sigma \in SL$ a security level. If $pc \vdash S$ and $\langle \mathfrak{G}, \mathfrak{L} \rangle \triangleright S \xrightarrow{\alpha} \langle \mathfrak{G}', \mathfrak{L}' \rangle \triangleright S'$ then $pc \vdash S'$.*

Subject reduction is a crucial ingredient to prove that well-typedness is a sufficient condition to ensure noninterference.

Theorem 3 (Soundness of security types). *Let Σ be a security policy, S a system of apps and $H_\sigma \stackrel{\text{def}}{=} \{\alpha \in \mathcal{A} \mid \Sigma(\alpha) \not\preceq \sigma\}$ the set of all possible high-level actions. If $\vdash S$ then*

$$\forall \mathfrak{G}, \mathfrak{G}' \in \mathbb{S}_\perp. \mathfrak{G} \equiv_\sigma \mathfrak{G}' \Rightarrow \langle \mathfrak{G}, \mathfrak{L}_\perp \rangle \triangleright S \approx_{H_\sigma} \langle \mathfrak{G}', \mathfrak{L}_\perp \rangle \triangleright S.$$

Proof. (Sketch) The proof proceeds by contradiction. Consider two configurations $\mathfrak{C} \stackrel{\text{def}}{=} \langle \mathfrak{G}, \mathfrak{L}_\perp \rangle \triangleright S$ and $\mathfrak{C}' \stackrel{\text{def}}{=} \langle \mathfrak{G}', \mathfrak{L}_\perp \rangle \triangleright S$ for some well-typed system S , and differing for two σ -equivalent global stores \mathfrak{G} and \mathfrak{G}' , respectively. If \mathfrak{C} and \mathfrak{C}' are not bisimilar then whenever we try to build up a termination-insensitive hiding bisimulation \mathcal{R} containing the pair $(\mathfrak{C}, \mathfrak{C}')$, the bisimulation game will stop in a pair of configurations $(\mathfrak{C}_i, \mathfrak{C}'_i)$, with \mathfrak{C}_i (*resp.*, \mathfrak{C}'_i) derivative of \mathfrak{C} (*resp.*, \mathfrak{C}'), such that \mathfrak{C}_i can perform some action α that cannot be (weakly) mimicked by \mathfrak{C}'_i (or vice versa). Since both \mathfrak{C}_i and \mathfrak{C}'_i are derivatives of well-typed configurations (actually well-typed systems), by subject reduction (Proposition 1) the corresponding systems are both well-typed.

Now, since our termination-insensitive hiding bisimilarity neglects high-level actions, by a case analysis on the distinguishing action α , the only possibility is that α denotes the update of a low-level service derived by the evaluation of some high-level service, the only ones on which the two global stores may differ. But this would lead to a security information

flow that is prevented by the type systems as both \mathcal{C}_i and \mathcal{C}'_i are well-typed. Thus, this action is not admissible.

As a consequence, since there are no distinguishing system actions α , the two original configurations \mathcal{C} and \mathcal{C}' must be bisimilar.

Full details of the proof can be found in the Appendix. \square

Example 8. Let us show that system S of Example 7 is cannot be typed in our type system. The process P running inside the applet Tw2Fb is typable: $L \vdash \Gamma\{P\}\Gamma$, for $\Gamma = [\text{tw} \mapsto H, \text{fb} \mapsto H, \text{ld} \mapsto L]$. As a consequence, the app is typable as well: $\vdash \text{Tw2Fb}[\text{tw}^R; \text{fb}^W \times P]$. On the other hand, the process Q running within the applet Fb2Ld is not typable. In fact, when building the derivation tree for the typing, we have to type the assignment $\text{ld} \leftarrow \text{fb}$ as $L \vdash \Gamma_1\{\text{ld} \leftarrow \text{fb}\}\Gamma_2$, for $\Gamma_1 = [\text{tw} \mapsto H, \text{fb} \mapsto H, \text{ld} \mapsto L]$ and $\Gamma_2 = [\text{tw} \mapsto H, \text{fb} \mapsto H, \text{ld} \mapsto H]$. But then, the subsequent update $\text{update}(\text{ld})$ cannot be typed, since the current typing of ld is H while the initial typing of ld is $\Sigma(\text{id}) = L$. Thus, the app $\text{Fb2Ld}[\text{fb}^R; \text{ld}^W \times Q]$ is not typable and, in turn, the parallel composition S is not typable.

We remark that our security type system did not face any permissiveness issues (i.e., false positives) for the apps considered in the paper. We expect our analysis to scale well and produce a minimal false-positive rate for user-automation IoT platforms like IFTTT and Stringify. In these platforms, the code consist of simple snippets matching the syntax of CaITApp closely [6]. For other IoT platforms like SmartThings, the code can be more complex (in fact, SmartThings apps are implemented in Groovy), hence our analysis would face classical challenges for type-based approaches due to complex language features, e.g., *aliasing*.

VI. RELATED WORK

Security and safety risks in the IoT domain have been the subject of a large array of research studies. We refer to recent surveys for an overview of the area [2], [5], [12]. Here, we compare our contributions with closely related works on security and safety analysis of IoT apps, information-flow control, and formal models for IoT.

Securing IoT apps: Recent research points out the security and safety risks arising in the context of IoT apps. Surbatovich et al. [36] study a dataset of 20K IFTTT applets and provide an empirical evaluation of potential secrecy and integrity violations, including violations due to cross-app interactions. Celik et al. [13], [14] propose static and dynamic enforcement mechanisms for unveiling cross-app interference vulnerabilities. Ding et al. [17] propose a framework that combines device physical channel analysis and static analysis to generate all potential interaction chains among applications in an IoT environment. They leverage natural language processing to identify services that have similar semantics, and propose a risk-based approach to classify the actual risks of the discovered interaction chains. Chi et al. [15] propose a systematic categorization of threats arising from unintentional or malicious interaction of apps in IoT platforms.

To detect cross-app interference, they use symbolic execution techniques to analyze the apps's implementation. Nguyen et al. [31] design IoTSan, a system that uses model checking to reveal cross-app interaction flows. All the above-mentioned works provide an excellent motivation for our foundational contributions. Our policy framework can be used to validate soundness and permissiveness of these verification techniques.

Another line of work focuses on enforcement mechanisms for checking security and safety of single IoT apps. Fernandes et al. [18] present FlowFence, an approach building secure IoT apps via information-flow tracking and controlled declassification. Celik et al. [11] leverage static taint tracking to identify sensitive data leaks in IoT apps. Bastys et al. [6], [7] identify new attack vectors in IFTTT applets and show that 30% of applets from their dataset can be subject to such attacks. As a countermeasure, they investigate static and dynamic information-flow tracking via security types. Fernandes et al. [19] propose the use of decentralization and fine-grained authentication tokens to limit privileges and prevent unauthorized actions. In contrast, our work targets security and safety issues in cross-app interactions, and it focuses on the formal underpinnings of these approaches.

Information-flow control: Several works propose information-flow control for enforcing confidentiality and integrity policies in emerging domains like IoT. We refer to a survey by Focardi and Gorrieri [20] for an overview on information-flow properties in process algebra. Our semantic condition of safe cross-app interaction draws inspiration from Focardi and Martinelli's notion of Generalized Non Deducibility on Composition (GNDC) [21], applied along the lines of [27]. Volpano and Smith [37] study a *flow-insensitive* type system for imperative languages. Because in our language the communication between services is handled via explicit *update* statements, a flow-insensitive type system would be too restrictive and reject more secure programs. Hunt and Sands [23] propose a flow-sensitive type system for an imperative language. Our work extends their type system to ensure security for a system of apps running concurrently. Similarly to our definition of termination-insensitive hiding bisimulation, Demange and Sands [16] propose a weakening of low bisimulation conditions to ignore leaks arising from program termination. In contrast, the execution of our app's payload affects the global store via a well-defined interface, i.e., listeners and update statements, which makes our systems of apps more amenable for enforcing security and safety properties.

There are a few approaches that carry out information-flow analysis on discrete/continuous models for cyber-physical systems. Akella et al. [1] proposed an approach to perform information flow analysis, including both trace-based analysis and automated analysis through process algebra specifications. This approach has been used to verify process algebra models of a gas pipeline system and a smart electric power grid system. Wang [38] propose Petri-net models to verify *nondeducibility security properties* of a natural gas pipeline system. More recently, Bohrer and Platzer [10] introduce dHL,

a hybrid logic for verifying cyber-physical hybrid-dynamic information flows, communicating information through both discrete computation and physical dynamics, ensuring security in presence of attackers that observe *continuously-changing values* in continuous time.

Formalizing IoT semantics: IoT semantics has been subject to several works aiming at capturing subtle IoT-specific notions like time and device state. Newcomb et al. [30] propose IOTA, a calculus for the domain of home automation. Based on the core formalism of IOTA, the authors develop an analysis for detecting whenever an event can trigger two conflicting actions, and an analysis for determining the root cause of (non)occurrence of an event. Lanese et al. [25] propose a calculus of mobile IoT devices interacting with the physical environment by means of sensors and actuators. The calculus does not allow any representation of the physical environment, while it is equipped with an end-user bisimilarity in which end-users may: (i) provide values to sensors, (ii) check actuators, and (iii) observe the mobility of smart devices. End-user bisimilarity is not preserved by parallel composition. Compositionality is recovered by strengthening its discriminating power. Lanotte and Merro [26] extend and generalize the work of [25] in a timed setting by providing a bisimulation-based semantic theory that is suitable for compositional reasoning. Bodei et al. [9] propose an untimed process calculus, IoT-LYSA, supporting a control flow analysis that safely approximates the abstract behavior of IoT systems. Essentially, they track how data spread from sensors to the logic of the network, and how physical data are manipulated. The calculus adopts asynchronous multi-party communication among nodes taking care of node proximity. The dynamics of the calculus is given in terms of a reduction relation. In [8], the same authors extend their work to infer *quantitative measures* to establish the cost of possibly security countermeasures, in terms of time and energy. In contrast, our calculus models constructs that are relevant for our security and safety analysis of cross-app interactions in IoT platforms, while ignoring details of the physical environment.

VII. CONCLUSION

IoT platforms empower users by connecting a wide array of otherwise unconnected services and devices. These platforms routinely execute IoT apps that have access to sensitive information of their users. Because different apps of a user may affect a common physical or logical environment, their interaction (even for benign apps) can cause severe security and safety risks for that user.

Motivated by this setting, we proposed a generic foundational framework for securing cross-app interactions. We presented an extensional condition that captures the essence of safe cross-app interactions, as well as implicit interactions and priorities. Moreover, we studied an extensional condition for confidentiality and integrity properties of a system of apps, and proposed a flow-sensitive security type system to enforce such condition.

Because of the simplicity of the execution model and the relatively-small code size, IoT apps offer a promising avenue for integrating formal techniques studied by the community to real-world products. In future work, we plan to implement our techniques for IoT platforms like Android Things and Node-RED, and investigate their impact on users' security and safety.

ACKNOWLEDGMENT

The authors would like to thank the anonymous reviewers for their valuable comments. The first author was partly funded by the Swedish Research Council (VR) under the project Joint-Force, and by the Swedish Foundation for Strategic Research (SSF) under the project TrustFull. The second author has been partially supported by the project "Dipartimenti di Eccellenza 2018-2022", funded by the Italian Ministry of Education, Universities and Research (MIUR). The third author has been supported by the Joint Project 2017 "Security Static Analysis for Android Things", jointly funded by the University of Verona and JuliaSoft Srl.

REFERENCES

- [1] R. Akella, H. Tang, and B. M. McMillin. Analysis of information flow security in cyber-physical systems. *International Journal of Critical Infrastructure Protection*, 3(3-4):157–173, 2010.
- [2] O. Alrawi, C. Lever, M. Antonakakis, and F. Monrose. SoK: Security evaluation of home-based iot deployments. In *IEEE S&P*, pages 208–226. IEEE Computer Society, 2019.
- [3] R. M. Amadio, I. Castellani, and D. Sangiorgi. On bisimulations for the asynchronous pi-calculus. *Theoretical Computer Science*, 195(2):291–324, 1998.
- [4] A. Askarov, S. Hunt, A. Sabelfeld, and D. Sands. Termination-insensitive noninterference leaks more than just a bit. In *ESORICS*, volume 5283 of *LNCS*, pages 333–348. Springer, 2008.
- [5] M. Balliu, I. Bastys, and A. Sabelfeld. Securing IoT Apps. *IEEE Security & Privacy Magazine*, September/October 2019.
- [6] I. Bastys, M. Balliu, and A. Sabelfeld. If This Then What? Controlling Flows in IoT Apps. In *ACM CCS*, pages 1102–1119. ACM, 2018.
- [7] I. Bastys, F. Piessens, and A. Sabelfeld. Tracking information flow via delayed output - Addressing privacy in IoT and emailing apps. In *NordSec*, volume 11252 of *LNCS*, pages 19–37. Springer, 2018.
- [8] C. Bodei, S. Chessa, and L. Galletta. Measuring security in IoT communications. *Theoretical Computer Science*, 764:100–124, 2019.
- [9] C. Bodei, P. Degano, G. Ferrari, and L. Galletta. Tracing where IoT data are collected and aggregated. *Logical Methods in Computer Science*, 13(3):1–38, 2017.
- [10] B. Bohrer and A. Platzer. A Hybrid, Dynamic Logic for Hybrid-Dynamic Information Flow. In *LICS*, pages 115–124. ACM, 2018.
- [11] Z. B. Celik, L. Babun, A. K. Sikder, H. Aksu, G. Tan, P. D. McDaniel, and A. S. Uluagac. Sensitive information tracking in commodity IoT. In *USENIX*, pages 1687–1704. USENIX Association, 2018.
- [12] Z. B. Celik, E. Fernandes, E. Pauley, G. Tan, and P. D. McDaniel. Program analysis of commodity IoT applications for security and privacy: Challenges and opportunities. *CoRR*, arXiv:1809.06962, 2018.
- [13] Z. B. Celik, P. McDaniel, and G. Tan. Soteria: Automated iot safety and security analysis. In *USENIX*, pages 147–158, 2018.
- [14] Z. B. Celik, G. Tan, and P. D. M. and. IoTGuard: Dynamic Enforcement of Security and Safety Policy in Commodity IoT. In *NDSS*. The Internet Society, 2019.
- [15] H. Chi, Q. Zeng, X. Du, and J. Yu. Cross-app interference threats in smart homes: Categorization, detection and handling. *CoRR*, abs/1808.02125, 2018.
- [16] D. Demange and D. Sands. All Secrets Great and Small. In *ESOP*, volume 5502 of *LNCS*, pages 207–221, 2009.
- [17] W. Ding and H. Hu. On the safety of IoT device physical interaction control. In *ACM CCS*, pages 832–846. ACM, 2018.
- [18] E. Fernandes, J. Paupore, A. Rahmati, D. Simonato, M. Conti, and A. Prakash. FlowFence: Practical Data Protection for Emerging IoT Application Frameworks. In *USENIX*, pages 531–548, 2016.

- [19] E. Fernandes, A. Rahmati, J. Jung, and A. Prakash. Decentralized Action Integrity for Trigger-Action IoT Platforms. In *NDSS*. The Internet Society, 2018.
- [20] R. Focardi and R. Gorrieri. Classification of security properties (Part I: Information flow). In *FOSAD*, pages 331–396. Springer-Verlag, 2001.
- [21] R. Focardi and F. Martinelli. A Uniform Approach for the Definition of Security Properties. In *FM*, volume 1708 of *LNCS*, pages 794–813. Springer, 1999.
- [22] K. Honda and M. Tokoro. An object calculus for asynchronous communications. In *ECOOP*, volume 512 of *LNCS*. Springer, 1991.
- [23] S. Hunt and D. Sands. On flow-sensitive security types. In *POPL*, pages 79–90. ACM, 2006.
- [24] IFTTT: IF This Then That. <https://ifttt.com>, 2019.
- [25] I. Lanese, L. Bedogni, and M. Di Felice. Internet of Things: a process calculus approach. In *SAC*, pages 1339–1346. ACM, 2013.
- [26] R. Lanotte and M. Merro. A semantic theory of the Internet of Things. *Information and Computation*, 259(1):72–101, 2018.
- [27] D. Macedonio and M. Merro. A semantic analysis of key management protocols for wireless sensor networks. *Science of Computer Programming*, 81:53–78, 2014.
- [28] M. Merro and D. Sangiorgi. On asynchrony in name-passing calculi. *Mathematical Structures in Computer Science*, 14(5):715–767, 2004.
- [29] Microsoft Flow. <https://flow.microsoft.com/en-us/>, 2019.
- [30] J. L. Newcomb, S. Chandra, J.-B. Jeannin, C. Schlesinger, and M. Sridharan. IOTA: A calculus for internet of things automation. In *International Symposium on New Ideas, New Paradigms, and Reflections on Programming and Software*. Onward!, pages 119–133, 2017.
- [31] D. T. Nguyen, C. Song, Z. Qian, S. V. Krishnamurthy, E. J. M. Colbert, and P. McDaniel. IoTSan: Fortifying the Safety of IoT Systems. In *CoNEXT*, pages 191–203. ACM, 2018.
- [32] Nest Thermostat. https://ifttt.com/services/nest_thermostat, 2019.
- [33] A. Sabelfeld and A. C. Myers. Language-based information-flow security. *IEEE Journal on Selected Areas in Communications*, 21(1):5–19, 2003.
- [34] Stringify. <https://www.stringify.com/>, 2019.
- [35] SmartThings. <https://ifttt.com/smartthings>, 2019.
- [36] M. Surbatovich, J. Aljuraidan, L. Bauer, A. Das, and L. Jia. Some recipes can do more than spoil your appetite: Analyzing the security and privacy risks of IFTTT recipes. In *WWW*, pages 1501–1510. ACM, 2017.
- [37] D. M. Volpano, C. E. Irvine, and G. Smith. A sound type system for secure flow analysis. *Journal of Computer Security*, 4(2/3):167–188, 1996.
- [38] J. Wang and H. Yu. Analysis of the Composition of Non-Deducibility in Cyber-Physical Systems. *Applied Mathematics & Information*, 8:3137–3143, 2014.

APPENDIX

A. Proofs

Proof of Theorem 1. It suffices to prove the first part of the theorem (for the second part, just swap S with R). Assume $S \not\leftrightarrow R$, i.e., $\text{act}(S) \cap \text{act}(R) = \emptyset = \text{act}(S) \cap \text{trg}(R)$, then we have to prove that for any global store \mathfrak{G} we have: $\langle \mathfrak{G}, \mathcal{L}_\perp \rangle \triangleright S \parallel R \approx_{H_S} R$, with $H_S = \text{up}\partial(S)$. The proof is by contradiction.

Suppose S and R be two syntactically noninteracting systems such that $\langle \mathfrak{G}, \mathcal{L}_\perp \rangle \triangleright S \parallel R \not\approx_{H_S} R$. By definition of bisimulation, this means that whenever we try to build up a hiding bisimulation \mathcal{R} containing the pair $(\mathcal{C}, \mathcal{C}')$, with $\mathcal{C} \stackrel{\text{def}}{=} \langle \mathfrak{G}, \mathcal{L}_\perp \rangle \triangleright S \parallel R$ and $\mathcal{C}' \stackrel{\text{def}}{=} \langle \mathfrak{G}, \mathcal{L}_\perp \rangle \triangleright R$, the bisimulation game will stop with a pair of configurations $(\mathcal{C}_i, \mathcal{C}'_i)$, with \mathcal{C}_i (resp., \mathcal{C}'_i) derivative of \mathcal{C} (resp., \mathcal{C}'), where \mathcal{C}_i can perform an action α that cannot be (weakly) mimicked by \mathcal{C}'_i (or vice versa).

Let proceed by case analysis on the action α that would distinguish the two configurations.

– $\alpha = \tau$. We notice that τ -actions cannot distinguish the two systems as we adopted a weak notion of bisimulation.

– $\alpha = x?v$. This action can be only derived by an application of rule (EnvChange) in Table II. However, as already pointed out, this action denotes a modification of the cloud made by the external observer. Thus, this action does not depend on the actual configuration and can always be performed by both configurations.

– $\alpha = \text{id}:x!v$. We have two sub-cases.

- id is an applet of the system S . In this case, $\alpha \in H_S$, and by definition of hiding bisimulation this action can always be mimicked by an arbitrary number (possibly 0) of τ -actions.

- id is an applet of the system R . In this case, $\alpha \notin H_S$. As α is the distinguishing action, it follows that the app id reaches different states in the two configurations \mathcal{C}_i and \mathcal{C}'_i leading to two possible situations: (i) the writing on x is possible in \mathcal{C}_i but not in \mathcal{C}'_i ; (ii) the writing on x is possible in both configurations but with different values. Since both $S \parallel R$ and R start in the same global store (the local store is not initialized in both cases), the system R could exhibit different behaviors if and only if it is affected by S . In particular, it means that S must modify a service that R listens on, or a service that R reads from the global store. Recall that there is no direct information passing between applets, so the only way for applets to interact is via the global store. However, we assumed that $\text{act}(S) \cap \text{trg}(R)$ is empty, hence this situation is not possible.

As it does not exist a distinguishing action α , it follows that the original configurations \mathcal{C} and \mathcal{C}' must be hiding bisimilar, with $H_S = \text{up}\partial(S)$. \square

Proof of Theorem 2. The proof goes along the same lines of that of Theorem 1. Indeed, the setting of the bisimulation is the same but the syntactic condition is stronger than the one of Theorem 1: $\text{act}(S) \subseteq \text{cl}\sigma(\Delta, \text{act}(S))$ and the relation $\overset{\Delta}{\leftrightarrow}$ (see Definition 8) is more restrictive than \leftrightarrow (see Definition 6). \square

Now, let us provide a technical lemma which is useful to prove the subject reduction property of the type system defined in Table III.

Lemma 1. *Let Σ be a security policy, P a process, and $\sigma \in SL$ a security level. If $\sigma \vdash \Gamma_1\{P\}\Gamma_2$ and $\langle \mathfrak{G}, \phi \rangle \triangleright P \xrightarrow{\lambda} \langle \mathfrak{G}', \mathcal{L}' \rangle \triangleright P'$ then there is Γ'_1 such that $\sigma \vdash \Gamma'_1\{P'\}\Gamma_2$.*

Proof. The proof is by rule induction on the transitions rules defining the semantics of processes (Table I). In that table, all transition rules are axioms (base cases of the induction) except for the rule (Seq) (the only inductive case). Let us proceed by case analysis on

$$\langle \mathfrak{G}, \phi \rangle \triangleright P \xrightarrow{\lambda} \langle \mathfrak{G}', \phi' \rangle \triangleright P'.$$

We examine the most significant cases.

- Rule (SetLocal) . In this case, we have $P = x \leftarrow e$, $P' = \text{skip}$, $\Gamma_1 \vdash e : \delta$, for some δ , and $\Gamma_2 = \Gamma_1[x \mapsto \delta \sqcup \sigma]$.

Then, setting $\Gamma'_1 = \Gamma_2$, we have $\sigma \vdash \Gamma'_1\{\text{skip}\}\Gamma_2$ by an application of the typing rule (Skip) .

- Rule (Update) . In this case, we have $P = \text{update}(x)$, $P' = \text{skip}$, $\Sigma(x) = \delta$, for some δ , $\Gamma_1 \vdash x : \delta'$, for some $\delta' \preceq \delta$ and $\Gamma_1 = \Gamma_2$. Then, setting $\Gamma'_1 = \Gamma_1 = \Gamma_2$, we have $\sigma \vdash \Gamma'_1\{\text{skip}\}\Gamma_2$ by an application of the typing rule (Skip) .
- Rule (IfTrue) . In this case, $P = \text{if } b \text{ then } \{P_1\} \text{ else } \{P_2\}$, $P' = P_1$ and $\Gamma_1 \vdash b : \delta$, for some δ . Then, by definition of the typing rule (IfElse) we have that $\sigma \sqcup \delta \vdash \Gamma'_1\{P_1\}\Gamma_2$, for some Γ'_1 . Finally, since $\sigma \preceq \sigma \sqcup \delta$, by an application of sub-typing we derive $\sigma \vdash \Gamma'_1\{P_1\}\Gamma_2$.
- Rule (Seq) . In this case, $P = P_1; P_2$, $\sigma \vdash \Gamma_1\{P_1\}\Gamma_3$, $\sigma \vdash \Gamma_3\{P_2\}\Gamma_2$, $\langle \mathcal{G}, \mathcal{L} \rangle \triangleright P_1 \xrightarrow{\lambda} \langle \mathcal{G}'', \mathcal{L}'' \rangle \triangleright P'_1$ and $P' = P'_1; P_2$. Since the derivation tree for $\langle \mathcal{G}, \mathcal{L} \rangle \triangleright P_1 \xrightarrow{\lambda} \langle \mathcal{G}'', \mathcal{L}'' \rangle \triangleright P'_1$ is smaller than the derivation tree for $\langle \mathcal{G}, \mathcal{L} \rangle \triangleright P \xrightarrow{\lambda} \langle \mathcal{G}', \mathcal{L}' \rangle \triangleright P'$, and $\sigma \vdash \Gamma_1\{P_1\}\Gamma_3$, for inductive hypothesis, we have that $\sigma \vdash \Gamma''_1\{P'_1\}\Gamma_3$, for some Γ''_1 . Thus, setting $\Gamma'_1 = \Gamma''_1$, since $\sigma \vdash \Gamma_3\{P_2\}\Gamma_2$, by an application of the typing rule (Seq) it follows that $\sigma \vdash \Gamma'_1\{P'_1; P_2\}\Gamma_2$.

□

Proof of Proposition 1. Given a security policy Σ , a system of apps S and a security level $\sigma \in SL$, we have to prove that if $\sigma \vdash S$ and $\langle \mathcal{G}, \mathcal{L} \rangle \triangleright S \xrightarrow{\alpha} \langle \mathcal{G}', \mathcal{L}' \rangle \triangleright S'$ then $\sigma \vdash S'$.

The proof is by rule induction on the transitions rules defining the semantics of systems (Table II). We proceed by case analysis on why

$$\langle \mathcal{G}, \mathcal{L} \rangle \triangleright S \xrightarrow{\alpha} \langle \mathcal{G}', \mathcal{L}' \rangle \triangleright S'.$$

- Rule (App) . In this case, we have $S = \text{id}[D \bowtie P]$, $\sigma \vdash \Gamma_1\{P\}\Gamma_2$, $\langle \mathcal{G}, \mathcal{L}(\text{id}) \rangle \triangleright P \xrightarrow{\alpha} \langle \mathcal{G}', \mathcal{L}'(\text{id}) \rangle \triangleright P'$ and $S' = \text{id}[D \bowtie P']$. By an application of Lemma 1, we derive that $\sigma \vdash \Gamma'_1\{P'\}\Gamma_2$, for some Γ'_1 . By an application of the typing rule (App) we derive $\sigma \vdash S'$.
- Rule (AppUpdate) . Similar to the previous case.
- Rule (EnvChange) . In this case, $S = S'$ and we trivially have $\sigma \vdash S'$.
- Rule (ParLeft) . In this case, we have $S = S_1 \parallel S_2$, $\langle \mathcal{G}, \mathcal{L} \rangle \triangleright S_1 \xrightarrow{\alpha} \langle \mathcal{G}'', \mathcal{L}'' \rangle \triangleright S'_1$, $\sigma \vdash S_1$ and $S' = S'_1 \parallel S_2$. Since the derivation tree for $\langle \mathcal{G}, \mathcal{L} \rangle \triangleright S_1 \xrightarrow{\alpha} \langle \mathcal{G}'', \mathcal{L}'' \rangle \triangleright S'_1$ is smaller than the derivation tree for $\langle \mathcal{G}, \mathcal{L} \rangle \triangleright S \xrightarrow{\alpha} \langle \mathcal{G}', \mathcal{L}' \rangle \triangleright S'$, and $\sigma \vdash S_1$, by inductive hypothesis it follows that $\sigma \vdash S'_1$. Thus, by an application of the typing rule (Par) it follows that $\sigma \vdash S'_1 \parallel S_2$.
- Rule (ParRight) . Similar to the previous case.

□

Proof of Theorem 3. Assume that $\vdash S$ and $\mathcal{G} \equiv_{\sigma} \mathcal{G}'$, then we have to prove that $\langle \mathcal{G}, \mathcal{L}_{\perp} \rangle \triangleright S \approx_{H_{\sigma}}^{\text{ti}} \langle \mathcal{G}', \mathcal{L}_{\perp} \rangle \triangleright S$, with $H_{\sigma} = \{\alpha \in \mathcal{A} \mid \Sigma(\alpha) \not\preceq \sigma\}$. The proof is by contradiction.

Suppose there exist two stores $\mathcal{G}, \mathcal{G}' \in \mathbb{S}_{\perp}$ such that $\mathcal{G} \equiv_{\sigma} \mathcal{G}'$ and $\langle \mathcal{G}, \mathcal{L}_{\perp} \rangle \triangleright S \not\approx_{H_{\sigma}}^{\text{ti}} \langle \mathcal{G}', \mathcal{L}_{\perp} \rangle \triangleright S$. This implies that whenever we try to build up a termination-insensitive hiding bisimulation \mathcal{R} containing the pair $(\mathcal{C}, \mathcal{C}')$, with $\mathcal{C} \stackrel{\text{def}}{=} \langle \mathcal{G}, \mathcal{L}_{\perp} \rangle \triangleright$

S and $\mathcal{C}' \stackrel{\text{def}}{=} \langle \mathcal{G}', \mathcal{L}_{\perp} \rangle \triangleright S$, the bisimulation game will stop in a pair of configurations $(\mathcal{C}_i, \mathcal{C}'_i)$, with \mathcal{C}_i (*resp.*, \mathcal{C}'_i) derivative of \mathcal{C} (*resp.*, \mathcal{C}'), where \mathcal{C}_i can perform an action α that cannot be (weakly) mimicked by \mathcal{C}'_i (or vice versa). Since both \mathcal{C}_i and \mathcal{C}'_i are derivatives of well-typed configurations (actually well-typed systems), by subject reduction (Proposition 1) the corresponding systems are both well-typed.

Let proceed by case a analysis on the action α that would distinguish the two configurations.

– $\alpha = \tau$. We notice that τ -actions cannot be a problem as we adopted a weak notion of bisimulation which can also mimic τ -actions.

– $\alpha = x?v$. This action can be only derived by an application of rule (EnvChange) in Table II. However, as already pointed out, this action denotes a modification of the cloud made by the external observer. Thus, this action does not depend on the actual configuration and can always be performed by both configurations.

– $\alpha = \text{id}:x!v$. We have two sub-cases.

• $\Sigma(x) \not\preceq \sigma$. That is α is a high-level action. In this case, $\alpha \in H_{\sigma}$, and by definition of our bisimulation this action can always be mimicked by an arbitrary number (possibly 0) of τ -actions.

• $\Sigma(x) \preceq \sigma$. That is α is a low-level action, or more precisely an action with security level smaller than or equal to σ . As α is the distinguish action, it follows that the app id reaches different states in the two configurations \mathcal{C}_i and \mathcal{C}'_i leading to two possible situations: (i) the low-level writing on x is possible in \mathcal{C}_i but not in \mathcal{C}'_i ; (ii) the low-level writing on x is possible in both configurations but with different values.

Let us consider the case (i) first. Since the initial global stores are σ -equivalent, $\mathcal{G} \equiv_{\sigma} \mathcal{G}'$, and potential updates of the global stores via the transition rule (EnvChange) will update the two stores consistently, it follows that $\mathcal{G}_i \equiv_{\sigma} \mathcal{G}'_i$, where \mathcal{G}_i (*resp.*, \mathcal{G}'_i) is the global store of \mathcal{C}_i (*resp.*, \mathcal{C}'_i). Thus, *one possibility* for the app id to reach two different states in the configurations \mathcal{C}_i and \mathcal{C}'_i , such that \mathcal{C}_i performs the update and \mathcal{C}'_i does not perform the update (or vice versa), is that it passed by a conditional statement that was evaluated differently in the two cases. Since $\mathcal{G}_i \equiv_{\sigma} \mathcal{G}'_i$, it follows that the guard of the conditional involved some high-level services which are the only ones that can differ in the two global stores. However, this situation is prevented by our type system as the typing rule (IfElse) ensures that we cannot have low-level updates in the two branches of the conditional, after the evaluation of a high-level guard (the high-level type is pushed into the *pc* of the branches and, eventually, to the low-level update). The *other possibility* is that there was a conditional statement that was evaluated differently in the two configurations: one branch diverges (without performing any low update) while the other does not. Then, after the conditional we have a low-level update. But in this case the two configurations are bisimilar, by definition of our termination-insensitive bisimulation.

Let us consider now the case (ii), *i.e.* when the low-level writing on x is possible in both configurations \mathcal{C}_i and \mathcal{C}'_i but with different values v and v' , respectively. This situation would be possible when the writing on x in the app id is preceded by an assignment $x \leftarrow e$ and the evaluations of e in the two configurations leads to v and v' , respectively, with $v \neq v'$. Since $\mathcal{G}_i \equiv_{\sigma} \mathcal{G}'_i$ it follows that the expression e must have a high-level security type, that is, $\Gamma \vdash e : \delta$, with $\sigma \prec \delta$.

A similar situation would be possible also if there was a conditional whose guard is typed at high-level and the service x is assigned in both branches of the conditional (with different values v and v'). Here, x would take different values even if the assignments involve only low-level expressions, since the guard has type greater than σ , similarly to what happens in case (i). However, our type system would type x as $\delta \succ \sigma$ and, via the typing rule (Update), it does not admit low-level writings on x as $\Sigma(x) \preceq \sigma \prec \delta = \Gamma(x)$. Thus, also this case is not admissible.

As it does not exist a distinguishing action α , it follows that the original configurations \mathcal{C} and \mathcal{C}' must be bisimilar, with $H_{\sigma} = \{\alpha \in \mathcal{A} \mid \Sigma(\alpha) \not\preceq \sigma\}$. \square

B. Further examples of IoT apps

In the following, we provide further examples of apps that can be modeled in our calculus CaITApp . These examples are inspired by existing real-world apps, also studied in the literature.

Example 9. Consider an app with the following specification. When the user enters a given geographical area the app sends an e-mail saying “entering area”. Similarly, when the user exits the given area, the app sends an e-mail saying “exiting area”. Suppose a services gps reporting the current position of the user in terms of gps coordinates: gps.Lat and gps.Long . Suppose also a service emailA , sending e-mails to a given address. The geographical area is given by a center, with coordinates centerLat and centerLong , and a radius called side . The area is the square with edges $(\text{centerLat} \pm \text{side}, \text{centerLong} \pm \text{side})$. The app is:

$$\text{Area}[\text{gps}^R; \text{emailA}^W \bowtie \text{fix } \mathbb{X} \bullet \text{listen}(\text{gps}); \text{pld9}]$$

where

$$\begin{aligned} \text{pld9} \stackrel{\text{def}}{=} & \text{if } (\text{ExitArea}) \text{ then } \{ \\ & \text{emailA} \leftarrow \text{goingOut}; \text{update}(\text{emailA}) \\ & \} \text{ else } \{ \\ & \text{if } (\text{EnterArea}) \text{ then } \{ \\ & \text{emailA} \leftarrow \text{goingIn}; \text{update}(\text{emailA}) \\ & \} \\ & \}; \\ & \text{gps.Lat} \leftarrow \text{read}(\text{gps.Lat}); \\ & \text{gps.Long} \leftarrow \text{read}(\text{gps.Long}); \\ & \mathbb{X} \end{aligned}$$

where ExitArea is a macro for the boolean test:

$$\begin{aligned} & \text{gps.Lat} - \text{centerLat} \leq \text{side} \wedge \\ & \text{gps.Long} - \text{centerLong} \leq \text{side} \wedge \\ & \text{read}(\text{gps.Lat}) - \text{centerLat} > \text{side} \wedge \\ & \text{read}(\text{gps.Long}) - \text{centerLong} > \text{side} \end{aligned}$$

and EnterArea is a macro for the boolean test:

$$\begin{aligned} & \text{gps.Lat} - \text{centerLat} > \text{side} \wedge \\ & \text{gps.Long} - \text{centerLong} > \text{side} \wedge \\ & \text{read}(\text{gps.Lat}) - \text{centerLat} \leq \text{side} \wedge \\ & \text{read}(\text{gps.Long}) - \text{centerLong} \leq \text{side} \end{aligned}$$

Example 10. Consider following app. When the lights are turned on the app should brew coffee and set the heater to 22 degrees. Suppose to have three services: light , managing the lights, coffeeM , managing the coffee machine, and heater , managing the heater. The app is formalized as follows:

$$\text{Welcome}[\text{light}^R; \text{coffeeM}^W; \text{heater}^W \bowtie \text{fix } \mathbb{X} \bullet \text{listen}(\text{light}); \text{P}_{10}]$$

where

$$\begin{aligned} \text{P}_{10} \stackrel{\text{def}}{=} & \text{lights} \leftarrow \text{read}(\text{lights}); \\ & \text{if } (\text{lights} = \text{On}) \text{ then } \{ \\ & \text{coffeeM} \leftarrow \text{doCoffee}; \\ & \text{heater} \leftarrow 22; \\ & \text{update}(\text{coffeeM}, \text{heater}) \\ & \}; \mathbb{X} \end{aligned}$$

Example 11. Imagine to have the following app. If a water leak is detected, then the app should shut off the main water supply valve. We assume two services: leakDetect , reporting whether there is a water leak or not, and mainValve , managing the main water valve. The app is formalized as follows:

$$\text{Leak}[\text{leakDetect}^R; \text{mainValve}^W \bowtie \text{fix } \mathbb{X} \bullet \text{listen}(\text{leakDetect}); \text{P}_{11}]$$

where

$$\begin{aligned} \text{P}_{11} \stackrel{\text{def}}{=} & \text{if } (\text{read}(\text{leakDetect}) = \text{yes}) \text{ then } \{ \\ & \text{mainValve} \leftarrow \text{Off}; \\ & \text{update}(\text{mainValve}) \\ & \}; \mathbb{X} \end{aligned}$$

Example 12. Imagine to have an Lamp1 , such that if the floor lamp is turned on and the home is in *sleep* mode, then the app should turn off the lamp after 5 minutes. We assume three services: flamp , managing the floor lamp, homeMode , reporting the home mode, and alarm , controlling the burglar alarm. The app is formalized as follows:

$$\text{Lamp1}[\text{flamp}^{\text{RW}}; \text{homeMode}^R \bowtie \text{fix } \mathbb{X} \bullet \text{listen}(\text{flamp}; \text{homeMode}); \text{P}_{12}]$$

where

$$P_{12} \stackrel{\text{def}}{=} \text{flamp} \leftarrow \text{read}(\text{flamp});$$

$$\quad \text{if } (\text{flamp.state} = \text{On} \wedge \text{homeMode} = \text{sleep}) \text{ then } \{$$

$$\quad \quad \text{flamp.ctrl} \leftarrow \text{OffDelay5};$$

$$\quad \quad \text{update}(\text{flamp})$$

$$\quad \}; \mathbb{X}$$

Example 13. Imagine to have a second app, Lamp2, such that if the floor lamp is turned on, and some object is moving, and it is midnight, and the lamp is turned on for more than 10 minutes, then the burglar alarm should be fired. We assume three services: flamp, managing the floor lamp, motion, reporting whether something is moving or not, and time, reporting the current time. The app is formalized as follows:

$$\text{Lamp2}[D_2 \bowtie \text{fix } \mathbb{X} \bullet \text{listen}(\text{flamp}; \text{motion}; \text{time}); P_{13}]$$

where

$$D_2 \stackrel{\text{def}}{=} \text{flamp}^{\text{RW}}; \text{motion}^{\text{R}}; \text{time}^{\text{R}}; \text{alarm}^{\text{W}}, \text{ and}$$

$$P_{13} \stackrel{\text{def}}{=} \text{flamp} \leftarrow \text{read}(\text{flamp});$$

$$\quad \text{motion} \leftarrow \text{read}(\text{motion});$$

$$\quad \text{time} \leftarrow \text{read}(\text{time});$$

$$\quad \text{if } (\text{cond}) \text{ then } \{$$

$$\quad \quad \text{if } (\text{flamp.curStateTime} \geq 10) \text{ then } \{$$

$$\quad \quad \quad \text{alarm} \leftarrow \text{On};$$

$$\quad \quad \quad \text{update}(\text{alarm})$$

$$\quad \quad \}$$

$$\quad \}$$

$$\quad \}; \mathbb{X}$$

where cond is the following boolean test:

$$\text{flamp.state} = \text{On} \wedge \text{motion} = \text{yes} \wedge \text{time.H} = 0 = \text{time.M.}$$

Example 14. Imagine to have an app Forward, forwarding messages from the e-mail address emailA to the e-mail address emailB. The app is formalized as follows:

$$\text{Forward}[\text{emailA}^{\text{R}}; \text{emailB}^{\text{W}} \bowtie \text{fix } \mathbb{X} \bullet \text{listen}(\text{emailA}); P_{14}]$$

where

$$P_{14} \stackrel{\text{def}}{=} \text{emailA} \leftarrow \text{read}(\text{emailA});$$

$$\quad \text{emailB} \leftarrow \text{emailA};$$

$$\quad \text{update}(\text{emailB}); \mathbb{X}$$

Let us now comment on potential interactions and on interferences between the apps defined above.

Semantic vs. syntax-based interactions: According to Definition 3, the apps Welcome and Leak do not interact with each other, since for every $\mathfrak{G} \in \mathbb{S}$ we have:

- $\langle \mathfrak{G}, \mathfrak{L}_\perp \rangle \triangleright \text{Welcome} \parallel \text{Leak} \approx_{H_w} \text{Leak}$
- $\langle \mathfrak{G}, \mathfrak{L}_\perp \rangle \triangleright \text{Welcome} \parallel \text{Leak} \approx_{H_l} \text{Welcome}$

where $H_w = \text{up}\delta(\text{Welcome})$ and $H_l = \text{up}\delta(\text{Leak})$.

However, if we add also the app SmokeAlarm, defined in Example 2, then we do have an interaction. In fact, there exists $\mathfrak{G} \in \mathbb{S}$ such that:

$$\langle \mathfrak{G}, \mathfrak{L}_\perp \rangle \triangleright \text{SmokeAlarm} \parallel \text{Welcome} \parallel \text{Leak} \not\approx_{H_s} \text{Welcome} \parallel \text{Leak}$$

where $H_s = \text{up}\delta(\text{SmokeAlarm})$. Here, the app SmokeAlarm may interact with the compound system Welcome \parallel Leak. Intuitively, when some smoke is detected the app SmokeAlarm turns the lights on. As a consequence, the app Welcome is triggered: the app brews a coffee and set the heater to 22 degrees.

According to Definition 6, this interaction can be captured syntactically as: SmokeAlarm \leftrightarrow Welcome \parallel Leak does not hold because $\text{act}(\text{SmokeAlarm}) = \{\text{alarm}, \text{lights}\}$, $\text{trg}(\text{Welcome} \parallel \text{Leak}) = \{\text{lights}, \text{leakDetect}\}$, and hence $\text{act}(\text{SmokeAlarm}) \cap \text{trg}(\text{Welcome} \parallel \text{Leak}) \neq \emptyset$.

Interactions under dependencies: According to Definition 3, the app Lamp1 is noninteracting with the app Lamp2, even if there is an obvious interplay between them. This lack in Definition 3 is due to the fact we cannot express that the action OffDelay5 turns flamp.state to Off after 5 minutes. This can be fixed by using a dependency policy $\Delta \stackrel{\text{def}}{=} \{(\text{flamp.ctrl}, \text{flamp.state})\}$. In fact, according to Definition 7, there exists $\mathfrak{G} \in \mathbb{S}$ such that:

$$\langle \mathfrak{G}, \mathfrak{L}_\perp \rangle \triangleright \text{Lamp1} \parallel \text{Lamp2} \not\approx_{H_1} \text{Lamp2}$$

with $H_1 \stackrel{\text{def}}{=} \{\text{Lamp1:flamp.ctrl!}v \in \mathcal{A} \mid v \in \text{Value}\}$. Intuitively, if the app Lamp1 turns off the lights then the app Lamp2 cannot be triggered. Again, this interaction can be captured syntactically, via Definition 8, as Lamp1 *stackrel*rel $\Delta \leftrightarrow$ Lamp2 does not hold because $\{\text{flamp.ctrl}, \text{flamp.state}\} \subseteq \text{cl}\sigma(\Delta, \text{act}(\text{Lamp1}))$, $\{\text{flamp.state}\} \subseteq \text{trg}(\text{Lamp2})$, and hence $\text{cl}\sigma(\Delta, \text{act}(\text{Lamp1})) \cap \text{trg}(\text{Lamp2}) \neq \emptyset$.

Interference and noninterference: Suppose that the email account A, associated to the address emailA, is confidential, i.e., $\Sigma(\text{emailA}) = \text{H}$, while the email account B, associated to the address emailB, is public, i.e., $\Sigma(\text{emailB}) = \text{L}$. For instance, emails associated to the account A can be read only by one, privileged, administrator while emails associated to the account B can be freely read by any user.

Clearly, the system $S \stackrel{\text{def}}{=} \text{Area} \parallel \text{Forward}$ is not secure as we have a confidentiality leak from emailA to emailB. Indeed, it is not hard to find two L-equivalent global stores \mathfrak{G} and \mathfrak{G}' such that $\langle \mathfrak{G}, \mathfrak{L}_\perp \rangle \triangleright S \not\approx_{H_l}^{\text{ti}} \langle \mathfrak{G}', \mathfrak{L}_\perp \rangle \triangleright S$. Again, we can capture this interference by syntactic means: the system S cannot be typed by our security type system because the process P_{14} running within the applet Forward is not typable. In fact, when building the derivation tree for the typing, we have to type the assignment $\text{emailB} \leftarrow \text{emailA}$ as $L \vdash \Gamma_1 \{ \text{emailB} \leftarrow \text{emailA} \} \Gamma_2$, for $\Gamma_1 = [\text{emailA} \mapsto \text{H}, \text{emailB} \mapsto \text{L}, \dots]$ and $\Gamma_2 = [\text{emailA} \mapsto \text{H}, \text{emailB} \mapsto \text{H}, \dots]$. But then, the subsequent command $\text{update}(\text{emailB})$ cannot be typed, since the current typing of emailB is H while the initial typing of emailB is $\Sigma(\text{emailB}) = \text{L}$. Thus, the app Forward is not typable and, in turn, the parallel composition S is not typable.