Information and Computation ••• (••••) •••–•••

Contents lists available at ScienceDirect

Information and Computation

www.elsevier.com/locate/yinco

A semantic theory of the Internet of Things ☆

Ruggero Lanotte [a], Massimo Merro [b,*]

[a] Dipartimento di Scienza e Alta Tecnologia, Università degli Studi dell'Insubria, Via Valleggio 11, 22100 Como, Italy
[b] Dipartimento di Informatica, Università degli Studi di Verona, Strada le Grazie 15, 37134 Verona, Italy

ARTICLE INFO

Article history:
Received 29 March 2017
Received in revised form 2 December 2017
Available online xxxx

Keywords:
Internet of Things
Process calculus
Operational semantics
Behavioural semantics
Bisimulation

ABSTRACT

We propose a process calculus for modelling and reasoning on systems in the *Internet of Things* paradigm. Our systems interact both with the physical environment, via *sensors* and *actuators*, and with *smart devices*, via short-range and Internet channels. The calculus is equipped with a standard notion of labelled *bisimilarity* which is proved to be a coinductive characterisation of a well-known contextual equivalence. We use our semantic proof-methods to prove run-time properties of a non-trivial case study as well as system equalities.

1. Introduction

In the *Internet of Things* (IoT) paradigm, *smart devices* equipped with embedded technology automatically collect information from shared resources (e.g. Internet accesses, physical devices, etc.) and aggregate them to provide new services to end users [2]. The "things" commonly deployed in IoT systems are: *RFID tags*, for unique identification, *sensors*, to detect physical changes in the environment, and *actuators*, to pass information to the environment. To provide proper communication capabilities, smart devices are organised in networks which are based on the standard communication protocols of the Internet framework.

The range of IoT applications is rapidly increasing and already covers several domains [3,2,4]: (i) environmental monitoring, (ii) healthcare, (iii) personal and social, (iv) security and surveillance, (v) smart environment (home, offices, cities), (vi) transportation and logistics (automotive).

The research on IoT is currently focusing on practical applications such as the development of enabling technologies [5], ad hoc architectures [6], semantic web technologies [7], and cloud computing [2]. However, as pointed out by Lanese et al. [8], there is a lack of research in formal methodologies to model the interactions among system components, and to verify the correctness of the network deployment before its implementation.

The main goal of the current paper is to propose a new process calculus for IoT systems which supports a clear semantic theory for specifying and reasoning on IoT applications. Devising a calculus for modelling a new paradigm requires understanding and distilling, in a clean algebraic setting, the basic features of the paradigm. In order to point out the main ingredients of the IoT paradigm, we use a small example within the smart environment domain.

☆ An extended abstract appeared in the proceedings of the *8th International Conference on Coordination Models and Languages* (COORDINATION 2016), volume 9686 of *Lecture Notes in Computer Science*, pp. 157–174, Springer, 2016 [1].

* Corresponding author.
E-mail address: massimo.merro@univr.it (M. Merro).

https://doi.org/10.1016/j.ic.2018.01.001
0890-5401/

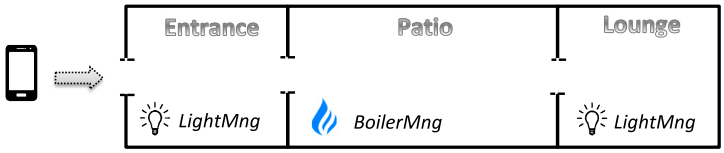

Fig. 1. A simple smart home.

Suppose a simple *smart home* (see Fig. 1) in which the user can (i) profit of her smartphone to remotely control the heating boiler of her house, and (ii) automatically turn on lights when entering a room. The house consists of an entrance and a lounge, separated by a patio. Entrance and lounge have their own lights (actuators) which are governed by different light manager processes, *LightMng*. The boiler is placed in the patio and it is governed by a boiler manager process, *BoilerMng*. This process senses the local temperature (via a sensor) and decides whether the boiler should be turned on/off, setting a proper actuator to signal the state of the boiler.

The smartphone executes two concurrent processes: *BoilerCtrl* and *LightCtrl*. The first one reads user's commands, submitted via the phone touchscreen (a sensor), and forwards them to the process *BoilerMng* of the house, via an Internet channel. Whereas, the process *LightCtrl* interacts with the processes *LightMng* of the house, via short-range wireless channels (e.g. Bluetooth, infrared, etc), to automatically turn on lights when the smartphone physically enters either the entrance or the lounge. The whole system is given by the parallel composition of the smartphone (a mobile device) and the smart home (a stationary entity).

On this kind of systems one may wish to prove interesting *run-time properties*. Think of a *fairness property* saying that the boiler will be eventually turned on/off whenever specific conditions are satisfied. Or *consistency properties*, saying, for instance, that the smartphone will never be in two rooms at the same time. Even more, one may be interested in understanding whether different implementations of our smart home have the same *observable behaviour*. Consider a variant of our smart home, where lights functionality depends on the GPS coordinates of the smartphone (localisation is a common feature of today smartphones). Intuitively, the smartphone could send its GPS position to a centralised light manager, *CLightMng* (possibly placed in the patio), via an Internet channel. The process *CLightMng* will then interact (via short-range channels) with the local light manager processes to turn on/off lights, depending on the current position of the smartphone. Here comes an interesting question: can these two implementations of the smart home, based on different light management mechanisms, be actually distinguished by an end user?

In the paper at hand we develop a fully abstract semantic theory for a process calculus of IoT systems, called `CaIT`. We provide a formal notion of when two systems in `CaIT` are indistinguishable, in all possible contexts, from the point of view of the end user. Formally, we adopt the approach of [9,10], often called *reduction (closed) barbed congruence*, which relies on two crucial concepts: a *reduction semantics* to describe system computations, and *basic observables* to represent what the environment can directly observe of a system. As IoT systems are essentially *cyber-physical systems* [11], they have at least two possible observables: the ability to transmit along channels, *logical observation*, and the capability to modify actuators, *physical observation*. In `CaIT`, we have adopted the second form of observable as our contextual equality remains invariant when adding logical observation. However, the right definition of physical observation is far from obvious as it has a non-trivial impact on the definition of the reduction semantics. Thus, observables and reduction semantics contain *key design choices* for the formal definition of `CaIT`.

Our calculus is equipped with two *labelled transition semantics* (LTSs) in the SOS style of Plotkin [12]: an *intensional* semantics and an *extensional* semantics. The adjective intensional is used to stress the fact that the actions here correspond to activities which can be performed by a system in isolation, without any interaction with the external environment. On the other hand, the extensional semantics focuses on those activities which require a contribution of the environment. Our extensional LTS builds on the intensional one, by introducing specific transitions for modelling all interactions with the environment. Here, we would like to point out that since our basic observation on systems does not involve the recording of the passage of time, this has to be taken into account extensionally.

We prove that the reduction semantics coincides with the intensional semantics (Harmony theorem), and that is satisfies some desirable time properties such as (a localised variant of) *time determinism*, *patience*, *maximal progress* and *well-timedness* [13]. However, the main result of the paper is that *weak bisimilarity* in the extensional LTS provides a *coinductive characterisation* of our contextual equivalence, reduction barbed congruence: two systems are related by some bisimulation in the extensional LTS if and only if they are reduction barbed congruent. Full abstraction results of this kind are in general hard to achieve. In our case, this result required a non-standard proof of the congruence theorem for the weak bisimilarity.

We finally show the effectiveness of our bisimulation proof-technique to deal with non-trivial systems. In particular, we provide a formal proof that two different implementations of the smart home mentioned before are bisimilar. Formal proofs of systems of such size are quite rare in the literature. Thus, in order to reduce the size of the bisimulation relation to be exhibited, we make an intensive use of *up-to expansion* proof-techniques [10].

Outline. Section 2 contains the calculus together with the reduction semantics, the contextual equivalence, and a discussion on design choices. Section 3 gives the details of our smart home example, and proves desirable run-time properties for it. Section 4 defines both intensional and extensional LTSs. In Section 5 we define bisimilarity for (networks of) IoT-systems,

Please cite this article in press as: R. Lanotte, M. Merro, A semantic theory of the Internet of Things, Inf. Comput. (2018), https://doi.org/10.1016/j.ic.2018.01.001

Table 1
Syntax.

Processes:

P, Q	::=	nil	termination
		$\rho.P$	intra-node activity
		$P \mid Q$	parallel composition
		$\lfloor \pi.P \rfloor Q$	communication with timeout
		$[b]P; Q$	conditional
		X	process variable
		$\mathsf{fix}\, X.P$	recursion

Networks:

M, N, O	::=	$\mathbf{0}$	empty network
		$n[\mathtt{I} \bowtie P]_h^\mu$	node/device
		$M \mid N$	network composition
		$(\nu c)M$	channel restriction

and prove the full abstraction result together with a number of non-trivial system equalities. Section 6 discusses related work, and concludes. Full details of the proofs can be found in the Appendix.

2. The calculus

In Table 1 we give the syntax of our *Calculus of the Internet of Things*, shortly CaIT, in a two-level structure: a lower one for *processes* and an upper one for *networks* of smart devices. We use letters n, m to denote *nodes/devices*, c, g for *channels*, h, k for (physical) *locations*, s, s' for *sensors*, a, a' for *actuators* and x, y, z for *variables*. Our *values*, ranged over by v and w, are constituted by basic values, such as booleans and integers, sensor and actuator values, and coordinates of physical locations.

A network is a pool of *distinct nodes* running in parallel. Nodes live in a physical world which can be divided in an enumerable set of physical locations. We assume a discrete notion of *distance* between two locations h and k, i.e. $\mathrm{d}(h,k) \in \mathbb{N}$. We write $\mathbf{0}$ to denote the empty network, while $M \mid N$ represents the parallel composition of two networks M and N. In $(\nu c)M$ channel c is private to the nodes of M. Each node is a term of the form $n[\mathtt{I} \bowtie P]_l^\mu$, where n is the device ID; $\mathtt{I}$ is the physical interface of n, represented as a partial mapping from sensor and actuator names to physical values; P is the process modelling the logics of n; l is the physical location of the device; $\mu \in \{\mathtt{s}, \mathtt{m}\}$ is a tag to distinguish between stationary and mobile nodes.

For security reasons, in a node $n[\mathtt{I} \bowtie P]_h^\mu$, sensors belonging to the physical interface $\mathtt{I}$ can be read only by the corresponding *controller process* P. Similarly, actuators in $\mathtt{I}$ can be modified only by P. No other devices can access the physical interface of n. P is a timed concurrent process which manages both the interaction with the physical interface $\mathtt{I}$ and channel communication with other devices. The communication paradigm is *point-to-point* via channels that may have different transmission ranges. We assume a global function rng() that given a channel c returns an element of $\mathbb{N} \cup \{-1, \infty\}$. Thus, a channel c can be used for: (i) *intra-node communications*, if $\mathrm{rng}(c) = -1$; (ii) *short-range inter-node communications* (such as Bluetooth) if $0 \leq \mathrm{rng}(c) < \infty$; (iii) *Internet communications*, if $\mathrm{rng}(c) = \infty$.

Our processes build on Hennessy and Regan's TPL [13] (basically, CCS with a discrete notion of time). We write $\rho.P$, with $\rho \in \{\sigma, @(x), s?(x), a!v\}$, to denote intra-node actions. The process $\sigma.P$ sleeps for one time unit. The process $@(x).P$ gets the current location of the enclosing node. Process $s?(x).P$ reads a value v from sensor s. Process $a!v.P$ writes the value v on the actuator a. We write $\lfloor \pi.P \rfloor Q$, with $\pi \in \{\overline{c}\langle v \rangle, c(x)\}$, to denote channel communication with timeout. This process can communicate along some channel c and, after that, it continues as P; otherwise, after one time unit, it evolves into Q. The process $[b]P; Q$ is the standard conditional construct, where b is a decidable guard. As in CCS, we assume that $[b]P; Q = P$ if $[\![b]\!] = \mathsf{true}$ (i.e. b evaluates to true), and $[b]P; Q = Q$ if $[\![b]\!] = \mathsf{false}$. In processes of the form $\sigma.Q$ and $\lfloor \pi.P \rfloor Q$ the occurrence of Q is said to be *time-guarded*. The process $\mathsf{fix}\, X.P$ denotes *time-guarded recursion*, as all occurrences of the process variable X may only occur time-guarded in P. In processes $\lfloor c(x).P \rfloor Q$, $s?(x).P$ and $@(x).P$ the variable x is said to be bound. Similarly, in process $\mathsf{fix}\, X.P$ the process variable X is bound. In the term $(\nu c)M$ the channel c is bound. This gives rise to the standard notions of *free/bound (process) variables*, *free/bound channels*, and *α-conversion*. A term is said to be *closed* if it does not contain free (process) variables, although it may contain free channels. We always work with closed networks: the absence of free variables is preserved at run-time. We write $T\{^v/_x\}$ for the substitution of the variable x with the value v in any expression T of our language. Similarly, $T\{^P/_X\}$ is the substitution of the process variable X with the process P in T.

The sensors embedded in a node can be of two kinds: *location-dependent* and *node-dependent*. The first ones sense data at the current location of the node, whereas the second ones sense data within the node, independently on the node's location. Thus, node-dependent sensor names are metavariables for sensors like *touchscreen@n* or *button@n*; whereas a sensor *temp@h*, for external temperature, is a typical example of location-dependent sensor. For simplicity, we use the same metavariables for both kinds of sensors. When necessary we will specify the type of sensor in use.

Actuator names are metavariables for actuators like *display@n* or *alarm@n*, where n is a node. As node names are unique so are actuator names: different nodes have different actuators. Thus, all actuators are basically node-dependent: location-dependent actuators would make little sense. Both actuator names and node-dependent sensor names are unique. This is not the case of location-dependent sensor names which may appear in different nodes.

The syntax given in Table 1 is a bit too permissive with respect to our intentions. We could rule out ill-formed networks with a simple type system. For the sake of simplicity, we prefer to provide the following definition.

Please cite this article in press as: R. Lanotte, M. Merro, A semantic theory of the Internet of Things, Inf. Comput. (2018), https://doi.org/10.1016/j.ic.2018.01.001

Table 2
Structural congruence.

Processes:

$P \mid \mathsf{nil} \equiv P$
$P \mid Q \equiv Q \mid P$
$(P \mid Q) \mid R \equiv P \mid (Q \mid R)$
$[b]P; Q \equiv P$ if $[\![b]\!] = \mathsf{true}$
$[b]P; Q \equiv Q$ if $[\![b]\!] = \mathsf{false}$
$\mathsf{fix}\, X.P \equiv P\{^{\mathsf{fix}\, X.P}/_X\}$

Networks:

$P \equiv Q$ implies $n[\mathtt{I} \bowtie P]^{\mu}_{h} \equiv n[\mathtt{I} \bowtie Q]^{\mu}_{h}$
$M|\mathbf{0} \equiv M$
$M|N \equiv N|M$
$(M|N)|O \equiv M|(N|O)$
$(\boldsymbol{\nu} c)\mathbf{0} \equiv \mathbf{0}$
$(\boldsymbol{\nu} c)(\boldsymbol{\nu} d)M \equiv (\boldsymbol{\nu} d)(\boldsymbol{\nu} c)M$
$(\boldsymbol{\nu} c)(M|N) \equiv M|(\boldsymbol{\nu} c)N$ if c not in M

Table 3
Reduction semantics.

$$(\mathrm{pos})\ \frac{-}{n[\mathtt{I} \bowtie @(x).P]^{\mu}_{h} \twoheadrightarrow_{\tau} n[\mathtt{I} \bowtie P\{^{h}/_{x}\}]^{\mu}_{h}} \qquad (\mathrm{sensread})\ \frac{\mathtt{I}(s) = v}{n[\mathtt{I} \bowtie s?(x).P]^{\mu}_{h} \twoheadrightarrow_{\tau} n[\mathtt{I} \bowtie P\{^{v}/_{x}\}]^{\mu}_{h}}$$

$$(\mathrm{actunchg})\ \frac{\mathtt{I}(a) = v}{n[\mathtt{I} \bowtie a!v.P]^{\mu}_{h} \twoheadrightarrow_{\tau} n[\mathtt{I} \bowtie P]^{\mu}_{h}} \qquad (\mathrm{actchg})\ \frac{\mathtt{I}(a) \neq v \quad \mathtt{I}' := \mathtt{I}[a \mapsto v]}{n[\mathtt{I} \bowtie a!v.P]^{\mu}_{h} \twoheadrightarrow_{a} n[\mathtt{I}' \bowtie P]^{\mu}_{h}}$$

$$(\mathrm{loccom})\ \frac{\mathrm{rng}(c) = -1}{n[\mathtt{I} \bowtie \lfloor \overline{c}\langle v\rangle.P \rfloor R \mid \lfloor c(x).Q \rfloor S]^{\mu}_{h} \twoheadrightarrow_{\tau} n[\mathtt{I} \bowtie P \mid Q\{^{v}/_{x}\}]^{\mu}_{h}}$$

$$(\mathrm{timestat})\ \frac{n[\mathtt{I} \bowtie \prod_i \lfloor \pi_i.P_i \rfloor Q_i \mid \prod_j \sigma.R_j]^{\mathtt{s}}_{h} \not\twoheadrightarrow_{\tau}}{n[\mathtt{I} \bowtie \prod_i \lfloor \pi_i.P_i \rfloor Q_i \mid \prod_j \sigma.R_j]^{\mathtt{s}}_{h} \twoheadrightarrow_{\sigma} n[\mathtt{I} \bowtie \prod_i Q_i \mid \prod_j R_j]^{\mathtt{s}}_{h}}$$

$$(\mathrm{timemob})\ \frac{n[\mathtt{I} \bowtie \prod_i \lfloor \pi_i.P_i \rfloor Q_i \mid \prod_j \sigma.R_j]^{\mathtt{m}}_{h} \not\twoheadrightarrow_{\tau} \quad \mathrm{d}(h,k) \leq \delta}{n[\mathtt{I} \bowtie \prod_i \lfloor \pi_i.P_i \rfloor Q_i \mid \prod_j \sigma.R_j]^{\mathtt{m}}_{h} \twoheadrightarrow_{\sigma} n[\mathtt{I} \bowtie \prod_i Q_i \mid \prod_j R_j]^{\mathtt{m}}_{k}}$$

$$(\mathrm{glbcom})\ \frac{\mathrm{d}(h,k) \leq \mathrm{rng}(c)}{n[\mathtt{I} \bowtie \lfloor \overline{c}\langle v\rangle.P \rfloor R]^{\mu_1}_{h} \mid m[\mathtt{J} \bowtie \lfloor c(x).Q \rfloor S]^{\mu_2}_{k} \twoheadrightarrow_{\tau} n[\mathtt{I} \bowtie P]^{\mu_1}_{h} \mid m[\mathtt{J} \bowtie Q\{^{v}/_{x}\}]^{\mu_2}_{k}}$$

$$(\mathrm{parp})\ \frac{\prod_i n_i[\mathtt{I}_i \bowtie P_i]^{\mu_i}_{h_i} \twoheadrightarrow_{\omega} \prod_i n_i[\mathtt{I}'_i \bowtie P'_i]^{\mu_i}_{h_i} \quad \omega \in \{\tau, a\}}{\prod_i n_i[\mathtt{I}_i \bowtie P_i|Q_i]^{\mu_i}_{h_i} \twoheadrightarrow_{\omega} \prod_i n_i[\mathtt{I}'_i \bowtie P'_i|Q_i]^{\mu_i}_{h_i}} \qquad (\mathrm{parn})\ \frac{M \twoheadrightarrow_{\omega} M' \quad \omega \in \{\tau, a\}}{M|N \twoheadrightarrow_{\omega} M'|N}$$

$$(\mathrm{timepar})\ \frac{M \twoheadrightarrow_{\sigma} M' \quad N \twoheadrightarrow_{\sigma} N' \quad M|N \not\twoheadrightarrow_{\tau}}{M|N \twoheadrightarrow_{\sigma} M'|N'} \qquad (\mathrm{timezero})\ \frac{-}{\mathbf{0} \twoheadrightarrow_{\sigma} \mathbf{0}}$$

$$(\mathrm{res})\ \frac{M \twoheadrightarrow_{\omega} N \quad \omega \in \{\tau, a, \sigma\}}{(\boldsymbol{\nu} c)M \twoheadrightarrow_{\omega} (\boldsymbol{\nu} c)N} \qquad (\mathrm{struct})\ \frac{M \equiv N \quad N \twoheadrightarrow_{\omega} N' \quad \omega \in \{\tau, a, \sigma\} \quad N' \equiv M'}{M \twoheadrightarrow_{\omega} M'}$$

Definition 2.1. A network M is said to be *well-formed if* (i) it does not contain two nodes with the same name; (ii) different nodes have different actuators; (iii) different nodes have different node-dependent sensors; (iv) for each $n[\mathtt{I} \bowtie P]^{\mu}_{h}$ in M, with a prefix $s?(x)$ (resp. $a!v$) in P, $\mathtt{I}(s)$ (resp. $\mathtt{I}(a)$) is defined; (v) for each $n[\mathtt{I} \bowtie P]^{\mu}_{h}$ in M with $\mathtt{I}(s)$ defined for some location-dependent sensor s, it holds that $\mu = \mathtt{s}$.

Condition (iv) requires that the physical devices (sensors and actuators) accessed by the controller processes must be defined in the corresponding physical interface of the node. Last condition implies that location-dependent sensors may be used only in stationary nodes. This restriction will be commented in Section 2.3. Hereafter, we will always work with *well-formed networks*. It is easy to show that well-formedness is preserved at runtime.

Finally, we assume a number of *notational conventions*. $\prod_{i\in I} M_i$ denotes the parallel composition of all M_i, for $i \in I$. We identify $\prod_{i\in I} M_i = \mathbf{0}$ and $\prod_{i\in I} P_i = \mathsf{nil}$, if $I = \emptyset$. Sometimes we write $\prod_i M_i$ when the index set I is not relevant. We write ρ instead of the process $\rho.\mathsf{nil}$. For $k \geq 0$, we write $\sigma^k.P$ as a shorthand for $\sigma.\sigma.\ldots.\sigma.P$, where prefix σ appears k consecutive times. Finally, we write $(\boldsymbol{\nu}\tilde{c})M$ as an abbreviation for $(\nu c_1)\ldots(\nu c_k)M$, for $\tilde{c} = c_1, \ldots, c_k$.

2.1. Reduction semantics

The dynamics of the calculus is given in terms of *reduction relations* over networks, as described in Table 3. As usual in process calculi, a reduction semantics relies on an auxiliary standard relation, $\equiv$, called *structural congruence*, which brings the participants of a potential interaction into contiguous positions. Formally, structural congruence is defined as the congruence induced by the axioms of Table 2 up to α-conversion.

As CaIT is a timed calculus, with a discrete notion of time, we will distinguish between instantaneous reductions, $M \twoheadrightarrow_{\mathrm{i}} N$, and timed reductions, $M \twoheadrightarrow_{\sigma} N$. Relation $\twoheadrightarrow_{\mathrm{i}}$ denotes activities which take place within one time interval, whereas $\twoheadrightarrow_{\sigma}$ represents the passage of one time unit. Our instantaneous reductions are of two kinds: those which involve the

change of the values associated to some actuator a, written $\rightarrow\!\triangleright_a$, and the others, written $\rightarrow\!\triangleright_\tau$. Intuitively, reductions of the form $M \rightarrow\!\triangleright_a N$ denote *watchpoints* which cannot be ignored by the physical environment (in Example 2.15, and more extensively at the end of Section 2.3, we explain why it is important to distinguish between $\rightarrow\!\triangleright_\tau$ and $\rightarrow\!\triangleright_a$). Thus, we define the instantaneous reduction relation $\rightarrow\!\triangleright_{\mathrm{i}} = \rightarrow\!\triangleright_\tau \cup \rightarrow\!\triangleright_a$, for any actuator a. We also define $\rightarrow\!\triangleright = \rightarrow\!\triangleright_\tau \cup \rightarrow\!\triangleright_\sigma$.

The first seven rules in Table 3 model intra-node activities. Rule (pos) serves to compute the current position of a node. Rule (sensread) represents the reading of the current data detected at some sensor s. Rules (actunchg) and (actchg) implement the writing of some data v on an actuator a, distinguishing whether the value of the actuator changes or not. Rule (loccom) models intra-node communications on a local channel c ($\mathrm{rng}(c) = -1$). Rule (timestat) models the passage of time within a stationary node. Notice that all untimed intra-node actions are considered urgent actions as they must occur before the next timed action. As an example, position detection is a time-dependent operation which cannot be delayed. Similar argument applies to sensor reading, actuator writing and channel communication. Rule (timemob) models the passage of time within a mobile node. This rule also serves to model *node mobility*. Mobile nodes can nondeterministically move from one physical location h to a (possibly different) location k, at the end of a time interval. Node mobility respects the following time discipline: in one time unit a node located at h can move to any location k such that $\mathrm{d}(h,k) \leq \delta$, for some fixed $\delta \in \mathbb{N}$ (obviously, it is possible to have $h = k$ and $\mathrm{d}(h,k) = 0$). For the sake of simplicity, we fix the same constant δ for all nodes of our systems. The premises of Rules (timestat) and (timemob) ensure that if a node can perform a timed reduction $\rightarrow\!\triangleright_\sigma$ then the same node cannot perform an instantaneous reduction $\rightarrow\!\triangleright_\tau$. Actually, due to the syntactic restrictions in the premises of both rules, that node cannot perform an instantaneous reduction $\rightarrow\!\triangleright_a$ either. This is formalised in Proposition 2.3.

Rule (glbcom) models inter-node communication along a global channel c ($\mathrm{rng}(c) \geq 0$). Intuitively, two different nodes can communicate via a common channel c if and only if they are within the transmission range of c. Rules (parp) and (parn) serve to propagate instantaneous reductions through parallel processes, and parallel networks, respectively. Rule (timepar) is for inter-node time synchronisation; the passage of time is allowed only if all instantaneous reductions have already fired. *Well-timedness* (Proposition 2.5) ensures the absence of infinite instantaneous traces which would prevent the passage of time. The remaining rules are standard.

We write $\rightarrow\!\triangleright_{\mathrm{i}}^{k}$ as a shorthand for k consecutive reductions $\rightarrow\!\triangleright_{\mathrm{i}}$; $\rightarrow\!\triangleright_{\mathrm{i}}^{*}$ is the reflexive and transitive closure of $\rightarrow\!\triangleright_{\mathrm{i}}$. Similar conventions apply to the reduction relation $\rightarrow\!\triangleright$.

Below we report a few standard time properties which hold in our calculus: *time determinism*, *maximal progress*, *patience* and *well-timedness*. In its standard formulation, *time determinism* says that a system reaches at most one new state by executing a reduction $\rightarrow\!\triangleright_\sigma$. However, by an application of Rule (timemob), our mobile nodes may change location when executing a reduction $\rightarrow\!\triangleright_\sigma$, thus we have a localised variant of time determinism.

Proposition 2.2 *(Localised time determinism). If* $M \rightarrow\!\triangleright_\sigma M'$ *and* $M \rightarrow\!\triangleright_\sigma M''$ *then* $M' \equiv \prod_{i\in I} n_i[\mathtt{I}_i \bowtie P_i]_{h_i}^{\mu_i}$ *and* $M'' \equiv \prod_{i\in I} n_i[\mathtt{I}_i \bowtie P_i]_{k_i}^{\mu_i}$, *with* $\mathrm{d}(h_i,k_i) \leq 2\delta$, *for all* $i \in I$.

According to [13], the *maximal progress* property says that processes communicate as soon as a possibility of communication arises. In CaIT, we generalise this property saying that instantaneous reductions cannot be delayed.

Proposition 2.3 *(Maximal progress). If* $M \rightarrow\!\triangleright_{\mathrm{i}} M'$ *then there is no* M'' *such that* $M \rightarrow\!\triangleright_\sigma M''$.

On the other hand, if no instantaneous reductions are possible then time is free to pass.

Proposition 2.4 *(Patience). If there is no* M' *such that* $M \rightarrow\!\triangleright_{\mathrm{i}} M'$ *then there is* N *such that* $M \rightarrow\!\triangleright_\sigma N$.

Finally, time passing cannot be prevented by infinite sequences of internal actions. Formally,

Proposition 2.5 *(Well-timedness). For any* M *there is a* $z \in \mathbb{N}$ *such that if* $M \rightarrow\!\triangleright_{\mathrm{i}}^{u} N$ *then* $u \leq z$.

The proof of Proposition 2.5 relies on time-guardedness of recursive processes.

2.2. Contextual behavioural equivalence

In this section, we provide a standard notion of contextual equivalence for our systems. Our touchstone equivalence is *reduction barbed congruence* [9,14], a standard contextually defined process equivalence. Intuitively, two systems are reduction barbed congruent if they have the same *basic observables* in all *contexts* and under all possible *computations*.

As already pointed out in the Introduction, the definition of reduction barbed congruence relies on two crucial concepts: a *reduction semantics* to describe system computations, and the *basic observable*, or barbs, which denotes what the environment can directly observe of a system.[1] So, the question now is: What are the "right" observables in our calculus? Due to

[1] See [10] for a comparison between this definition and the original definition of barbed congruence [14].

Please cite this article in press as: R. Lanotte, M. Merro, A semantic theory of the Internet of Things, Inf. Comput. (2018), https://doi.org/10.1016/j.ic.2018.01.001

the cyber-physical nature of our systems we could choose to observe either channel communications (logical observation) as in standard process calculi, or the capability to diffuse messages via actuators (physical observation). Actually, it turns out that logical observation in CaIT can be expressed in terms of physical one (see Section 2.3 for details) while the vice versa does not hold. So, we adopt as basic observables the capability to publish messages on actuators.

Definition 2.6 *(Barbs).* We write $M \downarrow_{a@h!v}$ if $M \equiv (\boldsymbol{\nu}\,\tilde{g})\big(n[\mathtt{I} \bowtie P]^{\mu}_{h} \mid M'\big)$ with $\mathtt{I}(a) = v$. We write $M \Downarrow_{a@h!v}$ if $M \rightarrow^{*} M'$ with $M' \downarrow_{a@h!v}$.

The reader may wonder why our barb reports the location and not the node of the actuator. We recall that actuator names are unique, so they somehow codify the name of their node. The location is then necessary because the environment is potentially aware of its position when observing an actuator: if on Monday at 6.00AM your smartphone rings to wake you up, then you may react differently depending whether you are at home or on holidays in the Bahamas!

Definition 2.7. A binary relation $\mathcal{R}$ over networks is *barb preserving* if $M\,\mathcal{R}\,N$ and $M \downarrow_{a@h!v}$ implies $N \Downarrow_{a@h!v}$.

Definition 2.8. A binary relation $\mathcal{R}$ over networks is *reduction closed* if whenever $M\,\mathcal{R}\,N$ the following conditions are satisfied:

- $M \rightarrow M'$ implies $N \rightarrow^{*} N'$ and $M'\,\mathcal{R}\,N'$;
- $M \rightarrow_{a} M'$ implies $N \rightarrow^{*} \rightarrow_{a} \rightarrow^{*} N'$ and $M'\,\mathcal{R}\,N'$.

Here, we require reduction closure of both $\rightarrow$ and $\rightarrow_{a}$, for any a. This is a *crucial design decision* in CaIT (see Example 2.15 and Section 2.3 for details).

In order to model sensor updates made by the physical environment on a sensor s in a given location h, we define an operator $[s@h \mapsto v]$ on networks.

Definition 2.9. Given a location h, a sensor s, and a value v in the domain of s, we define:

$$
\begin{array}{rcl}
n[\mathtt{I} \bowtie P]^{\mu}_{h}[s@h \mapsto v] & \stackrel{\text{def}}{=} & n[\mathtt{I}[s \mapsto v] \bowtie P]^{\mu}_{h}, \text{ if } \mathtt{I}(s) \text{ defined} \\
n[\mathtt{I} \bowtie P]^{\mu}_{k}[s@h \mapsto v] & \stackrel{\text{def}}{=} & n[\mathtt{I} \bowtie P]^{\mu}_{k}, \text{ if } \mathtt{I}(s) \text{ undefined or } h \neq k \\
(M|N)[s@h \mapsto v] & \stackrel{\text{def}}{=} & M[s@h \mapsto v] \mid N[s@h \mapsto v] \\
\big((\boldsymbol{\nu} c)M\big)[s@h \mapsto v] & \stackrel{\text{def}}{=} & (\boldsymbol{\nu} c)\big(M[s@h \mapsto v]\big) \\
\mathbf{0}[s@h \mapsto v] & \stackrel{\text{def}}{=} & \mathbf{0}.
\end{array}
$$

As for barbs, the reader may wonder why when updating a sensor we use its location, also for node-dependent sensors. This is because when changing a node-dependent sensor (e.g. touching a touchscreen of a smartphone) the environment is in general aware of its position.

Definition 2.10. A binary relation $\mathcal{R}$ over networks is *contextual* if $M\,\mathcal{R}\,N$ implies that

- for all networks O, $M|O\ \mathcal{R}\ N|O$;
- for all channels c, $(\boldsymbol{\nu} c)M\ \mathcal{R}\ (\boldsymbol{\nu} c)N$;
- for all sensors s, locations h, and values v in the domain of s, $M[s@h \mapsto v]\,\mathcal{R}\,N[s@h \mapsto v]$.

The first two clauses require closure under *logical contexts* (parallel systems), while the last clause involves *physical contexts*, which can *nondeterministically update* sensor values.

Finally, everything is in place to define our touchstone contextual behavioural equality.

Definition 2.11. *Reduction barbed congruence*, written $\cong$, is the largest symmetric relation over networks which is reduction closed, barb preserving and contextual.

Remark 2.12. Obviously, if $M \cong N$ then M and N will be equivalent in any setting where sensor updates are governed by specific physical laws. This is because physical contexts that can affect sensor values (according to some physical law) are definitely fewer than those which can change sensors nondeterministically.

We recall that the reduction relation $\rightarrow$ ignores the passage of time, and therefore the reader might suspect that our reduction barbed congruence is impervious to the precise timing of activities. We will show that this is not the case.

Please cite this article in press as: R. Lanotte, M. Merro, A semantic theory of the Internet of Things, Inf. Comput. (2018), https://doi.org/10.1016/j.ic.2018.01.001

Example 2.13. Let M and N be two networks such that $M = n[\emptyset \bowtie \sigma.\lfloor \overline{c}\langle v\rangle.\mathrm{nil}\rfloor \mathrm{nil}]_h^{\mathrm{s}}$ and $N = n[\emptyset \bowtie \lfloor \overline{c}\langle v\rangle.\mathrm{nil}\rfloor \mathrm{nil}]_h^{\mathrm{s}}$, with $\mathrm{rng}(c) = \infty$. It is easy to see that $M \rightarrow_\sigma N$. As the reduction relation $\rightarrow$ does not distinguish instantaneous reductions from timed ones, one may think that networks M and N are reduction barbed congruent, and that a prompt transmission along channel c is equivalent to the same transmission delayed of one time unit. However, let us consider the test $T = test[\mathtt{J} \bowtie \sigma.a!1.\lfloor c(x).a!0\rfloor \mathrm{nil}]_l^{\mathrm{s}}$, with $\mathtt{J}(a) = 0$, for some (fresh) actuator a. Our claim is that test T can distinguish the two networks, and thus $M \not\cong N$. In fact, if $M|T \rightarrow \rightarrow_a O = n[\emptyset \bowtie \lfloor \overline{c}\langle v\rangle.\mathrm{nil}\rfloor \mathrm{nil}]_h^{\mathrm{s}} \mid test[\mathtt{J}' \bowtie \lfloor c(x).a!0\rfloor \mathrm{nil}]_l^{\mathrm{s}}$, with $\mathtt{J}'(a) = 1$, then there is no O' such that $N|T \rightarrow^* \rightarrow_a \rightarrow^* O'$ with $O \cong O'$. This is because O can perform a reduction sequence $\rightarrow \rightarrow_a$ that cannot be matched by any O'.

Behind this example there is the general principle that reduction barbed congruence is sensitive to the passage of time.

Proposition 2.14. *If $M \cong N$ and $M \rightarrow_\sigma M'$ then there is N' such that $N \rightarrow_\tau^* \rightarrow_\sigma \rightarrow_\tau^* N'$ and $M' \cong N'$.*

Proof. Suppose $M \rightarrow_\sigma M'$. Consider the test node $T = test[\mathtt{J} \bowtie \sigma.a!1.a!0]_k^{\mathrm{s}}$ such that both networks $M|T$ and $N|T$ are well-formed, and $\mathtt{J}(a) = 0$. By construction, the presence of a barb $\Downarrow_{a@k!1}$ in a derivative of one of those systems implies that exactly one timed reduction $\rightarrow_\sigma$ has been inferred in the derivation.

Since $M \rightarrow_\sigma M'$ it follows that $M|T \rightarrow_\sigma \rightarrow_a M'|T'$, with $T' = test[\mathtt{J}[a \mapsto 1] \bowtie a!0]_k^{\mathrm{s}}$ and $M'|T' \downarrow_{a@k!1}$. As $M \cong N$ and $\cong$ is contextual, the reduction sequence above must be mimicked by $N|T$, that is $N|T \rightarrow^* \rightarrow_a \rightarrow^* \hat{N}$, with $M'|T' \cong \hat{N}$. As a consequence, $\hat{N} \Downarrow_{a@k!1}$. This implies that exactly one timed reduction has been inferred in the reduction sequence $N|T \rightarrow^* \rightarrow_a \rightarrow^* \hat{N}$. As $M|T$ and $N|T$ are well-formed networks, the actuator a can appear neither in M nor in N. So, the above reduction sequence can be decomposed as follows:

$$N|T \rightarrow^* N'|T \rightarrow_a N'|T' \rightarrow^* N''|T' = \hat{N}$$

with $N \rightarrow_\tau^* \rightarrow_\sigma \rightarrow_\tau^* N''$. From $M'|T' \cong N''|T'$ it is easy to derive $M' \cong N''$ (for details see Lemma 5.7 in Section 5). □

Now, we provide some insights into the design decision of having two different reduction relations $\rightarrow_\tau$ and $\rightarrow_a$.

Example 2.15. Let $M = n[\mathtt{I} \bowtie a!1|a!0.a!1]_h^{\mu}$ and $N = n[\mathtt{I} \bowtie a!1.a!0.a!1]_h^{\mu}$, with $\mathtt{I}(a) = 0$ and undefined otherwise. Then, within one time unit, M may display on the actuator a either the sequence of values 01 or the sequence 0101, while N can only display the sequence 0101. As a consequence, from the point of view of the physical environment, the observable behaviours of M and N are clearly different. In the following we show how $\cong$ can observe that difference. We recall that the relation $\cong$ is reduction closed. Now, if $M \rightarrow_\tau \rightarrow_a M' = n[\mathtt{J} \bowtie a!1]_h^{\mu}$, with $\mathtt{J}(a) = 1$, the only possible reply of N respecting reduction closure is $N \rightarrow^* \rightarrow_a N' = n[\mathtt{J} \bowtie a!0.a!1]_h^{\mu}$. However, it is evident that $M' \not\cong N'$ because N' can turn the actuator a to 0 while M' cannot. Thus, $M \not\cong N$.

Had we merged the relation $\rightarrow_a$ with $\rightarrow_\tau$ then we would have $M \cong N$ because the capability to observe messages on actuators, given by the barb, would not be enough to observe changes on actuators within one time interval.

2.3. Design choices

In this section we provide some insights into the design choices that have been followed in the definition of CaIT.

CaIT is a *value-passing* rather than a *name-passing* calculus, à la π-calculus [10]; the theory of CaIT can be easily adapted to deal with the transmission of channel names at the price of adding the standard burden of scope extrusion of names. Furthermore, as both actuators and sensors can only be managed inside their nodes, it would make little sense to transmit their names along channels. For simplicity, in CaIT we adopt a *point-to-point communication* via communication channels. Broadcast communications could be easily introduced along the lines of Cerone, Hennessy and Merro's timed calculus CCCP [15], without affecting the main theoretical results.

CaIT is a timed process calculus with a discrete notion of time. The time model we adopt in CaIT is known as the *fictitious clock approach* (see e.g. [13]): a global clock is supposed to be updated whenever all nodes agree on this, by globally synchronising on a special timing action σ. Thus, time synchronisation relies on some clock synchronisation protocol for mobile wireless systems [16]. However, our notion of time interval is different from that adopted in synchronous languages [17–19] where the environment injects events at the beginning of a time interval and collects them at the end. In the synchronous approach, events happening during a time interval are not ordered while in our calculus we want to maintain the causality among actions, typical of process calculi.

We already said that IoT systems are essentially *cyber-physical systems* [11]. In cyber-physical systems, sensor changes are usually modelled either using continuous models, such as differential equations, or through discrete models, such as difference equations.[2] However, in this paper we aim at providing a behavioural semantics for IoT applications from the

[2] Difference equations relate to differential equations as discrete math relate to continuous math.

Please cite this article in press as: R. Lanotte, M. Merro, A semantic theory of the Internet of Things, Inf. Comput. (2018), https://doi.org/10.1016/j.ic.2018.01.001

point of the view of the end user. And the end user cannot directly observe changes on the sensors of an IoT application: she can only observe the effects of those changes via actuators and communication channels. Thus, in CaIT we do not represent sensor changes via specific models, but we rather abstract on them by supporting *nondeterministic sensor updates* (see Definition 2.9 and Definition 2.10). Actually, as pointed out in Remark 2.12, behavioural equalities derived in our setting remain valid when adopting any specific model for sensor updates.

Another design decision in our language regards the possibility to change the value associated to sensors and actuators more than once within the same time interval. At first sight this choice may appear weird as certain actuators are physical devices that may require some time to turn on. On the other hand, other actuators, such as lights or displays, may have very quick reactions. A similar argument applies to sensors. In this respect our calculus does not enforce a synchronisation of physical events as for logical signals in synchronous languages. In fact, actuator changes are under nodes' control: the process running within a node decides when changing the value exposed on an actuator of that node. Thus, if the actuator of a node models a slow device then it is under the responsibility of the process running at that node to change the actuator with a proper delay. Similarly, sensors should be read only when this makes sense. For instance, a temperature sensor should be read only when the temperature is known to be stable.

Let us now discuss on node mobility. The reader may wonder why CaIT does not provide a process for node mobility, as in Mobile Ambients [20]. Notice that, unlike Mobile Ambients, our nodes do not represent mobile computations within an Internet domain. Instead, they represent smart devices which do not decide where to move to: an external agent moves them. We also decided to allow node mobility only at the end of time intervals. This is because both intra-node and inter-node logical operations, such as channel communications, can be considered significantly faster than physical movements of devices. For instance, consider a transmitter that moves at 20 m/s and that transmits a 2000-byte frame over a channel having a 2 megabit/s bandwidth. The actual transmission would take about 0.008 s; during that time, the transmitter moves only about 16 cm away. In other words, we can assume that the nodes are stationary when transmitting and receiving, and may change their location only while they are idle. However, to avoid uncontrolled movements of nodes we decided to fix for all of them the same bound δ, representing the maximum distance a node can travel within one time unit. There would not be problems in allowing different δ for different nodes. Finally, for the sake of simplicity, in the last constraint of Definition 2.1 we impose that location-dependent sensors can only occur in stationary nodes. This allows us to have a local, rather than a global, representation of those sensors. Notice that mobile location-dependent sensors would have the same technical challenges of *mobile wireless sensor networks* [21].

Another issue is about a proper representation of network topology. A tree-structure topology, as in Mobile Ambients, would be desirable to impose that a device cannot be in two mutually exclusive places at the same time. This desirable property cannot be expressed in [8], where links between nodes can be added and removed nondeterministically. However, a tree-structured topology would imply an higher-order bisimulation (for details see [22]); while in the current paper we look for a simple (first-order) bisimulation proof-technique which could be easily mechanised.

Finally, we would like to explain our choice about barbs. As already said in the previous section there are other possible definitions of barb. For instance, one could choose to observe the capability to transmit along a channel c, by defining $M \downarrow_{\tilde{c}@h}$ if $M \equiv (\boldsymbol{\nu}\tilde{g})\big(n[\mathtt{I} \bowtie \lfloor \overline{c}\langle v\rangle.P \rfloor P' \mid Q]^{\mu}_{k} \mid M'\big)$, with $c \notin \tilde{g}$ and $\mathrm{d}(h,k) \leq \mathrm{rng}(c)$. However, if you consider the system $S = (\boldsymbol{\nu}c)(M | m[\mathtt{J} \bowtie \lfloor c(x).a!1 \rfloor \mathsf{nil}]^{\mu}_{h})$, with $\mathtt{J}(a) = 0$, for some appropriate m, then it is easy to show that $M \downarrow_{\tilde{c}@h}$ if and only if $S \rightarrow\!\!\triangleright\!\!\rightarrow_{a} S'$ with $S' \downarrow_{a@h!1}$. Thus, the barb on channels can always be reformulated in terms of our barb. The vice versa is not possible. The reader may also wonder whether it is possible to turn the reduction $\rightarrow\!\!\triangleright_{a}$ into $\rightarrow\!\!\triangleright_{\tau}$ by introducing, at the same time, some special barb which would be capable to observe actuators changes. For instance, something like $M \downarrow_{a@h!v.w}$ if $M \equiv (\boldsymbol{\nu}\tilde{g})\big(n[\mathtt{I} \bowtie a!w.P | Q]^{\mu}_{h} \mid M'\big)$, with $\mathtt{I}(a) = v$ and $v \neq w$. It should be easy to see that this extra barb would not help in distinguishing the terms proposed in Example 2.15. Actually, here there is something deeper that needs to be spelled out. In process calculi, the term β of a barb $\downarrow_{\beta}$ is a concise encoding of a context C_{β} expressible in the calculus and capable to observe the barb $\downarrow_{\beta}$. However, our barb $\downarrow_{a@h!v}$ does not have such a corresponding *physical context* in our language. For instance, in CaIT we do not represent the "eyes of a person" looking at the values appearing to some display. Technically speaking, we do not have terms of the form $a?(x).P$ that could be used by the physical environment to read values on the actuator a. This is because such terms would not be part of an IoT system. The lack of this physical context, together with the persistent nature of actuators' state, explains why our barb $\downarrow_{a@h!v}$ must work together with the reduction relation $\rightarrow\!\!\triangleright_{a}$ to provide the desired distinguishing power of $\cong$. On the other hand, the decision of not including $\rightarrow\!\!\triangleright_{a}$ as part of $\rightarrow\!\!\triangleright$ gives to $\cong$ enough distinguishing power to observe strong preservation of barbs.

Proposition 2.16. *If $M \cong N$ and $M \downarrow_{a@h!v}$ then $N \downarrow_{a@h!v}$.*

Proof. We recall that $\rightarrow\!\!\triangleright \stackrel{\text{def}}{=} \rightarrow\!\!\triangleright_{\tau} \cup \rightarrow\!\!\triangleright_{\sigma}$. Suppose $M \downarrow_{a@h!v}$. As $\cong$ is barb preserving it follows that $N \Downarrow_{a@h!v}$, namely, there is N' such that $N \rightarrow\!\!\triangleright^{*} N'$ with $N' \downarrow_{a@h!v}$. However, both reduction relations $\rightarrow\!\!\triangleright_{\tau}$ and $\rightarrow\!\!\triangleright_{\sigma}$ do not modify actuator values. As a consequence, this holds also for $\rightarrow\!\!\triangleright$. Thus, it follows that $N \downarrow_{a@h!v}$. □

3. Case study: a smart home

In this section, we model the simple smart home discussed in the Introduction, and represented in Fig. 1. The house consists of an entrance and a lounge, separated by a patio. It spans over 4 contiguous physical locations *loci*, for $i \in$

Table 4
A smart home in CaIT.

$$
\begin{array}{rcl}
Sys & \stackrel{\text{def}}{=} & Phone \mid Home \\
Phone & \stackrel{\text{def}}{=} & n_P[\mathtt{I}_P \bowtie BoilerCtrl \mid LightCtrl]^{\mathrm{m}}_{out} \\
Home & \stackrel{\text{def}}{=} & LM_1 \mid LM_2 \mid BM \\
LM_1 & \stackrel{\text{def}}{=} & n_1[\mathtt{I}_1 \bowtie LightMng_1]^{\mathrm{s}}_{loc1} \\
LM_2 & \stackrel{\text{def}}{=} & n_2[\mathtt{I}_2 \bowtie LightMng_2]^{\mathrm{s}}_{loc4} \\
BM & \stackrel{\text{def}}{=} & n_B[\mathtt{I}_B \bowtie BoilerMng]^{\mathrm{s}}_{loc2} \\
BoilerCtrl & \stackrel{\text{def}}{=} & \mathsf{fix}\, X.mode?(z).\lfloor \overline{b}\langle z\rangle.\sigma.X \rfloor X \\
LightCtrl & \stackrel{\text{def}}{=} & \prod_{j=1}^{2} \mathsf{fix}\, X.\lfloor \overline{c_j}\langle \mathsf{on}\rangle.\sigma.X \rfloor X \\
LightMng_j & \stackrel{\text{def}}{=} & \mathsf{fix}\, X.\lfloor c_j(x).light_j!\mathsf{on}.\sigma.X \rfloor light_j!\mathsf{off}.X \quad \text{for } j \in \{1,2\} \\
BoilerMng & \stackrel{\text{def}}{=} & \mathsf{fix}\, X.\lfloor b(x).[x = \mathsf{man}]\, boiler!\mathsf{on}.\sigma.BoilerManual; TempCtrl \rfloor TempCtrl \\
BoilerManual & \stackrel{\text{def}}{=} & \mathsf{fix}\, Y.b(y).[y = \mathsf{auto}]X; \sigma.Y \\
TempCtrl & \stackrel{\text{def}}{=} & temp?(t).[t < \Theta]\, boiler!\mathsf{on}.\sigma.X; boiler!\mathsf{off}.\sigma.X
\end{array}
$$

$\{1,2,3,4\}$, such that $\mathrm{d}(loci, locj) = |i - j|$. The entrance is in $loc1$, the patio spans from $loc2$ to $loc3$, and the lounge is in $loc4$. The house can only be accessed via its entrance. Entrance and lounge have their own lights (actuators) which are governed by different light manager processes, *LightMng*. The boiler is in the patio and is governed by a boiler manager process, *BoilerMng*. This process senses the local temperature (via a sensor) and decides whether the boiler should be turned on/off, setting a proper actuator to signal the state of the boiler. The smartphone executes two concurrent processes: *BoilerCtrl* and *LightCtr*. The first one reads user's commands for the boiler, submitted via the phone touchscreen (a sensor), and forwards them to the process *BoilerMng*, via an Internet channel. Whereas, the process *LightCtrl* interacts with the processes *LightMng*, via short-range wireless channels (e.g. Bluetooth, infrared, etc), to automatically turn on lights when the smartphone physically enters either the entrance or the lounge.

Table 4 provides a detailed formalisation of our smart home in CaIT. The whole system, *Sys*, is given by the parallel composition of the smartphone *Phone* (a mobile device) and the smart home *Home* (a stationary entity). The smartphone is represented as a mobile node, with $\delta = 1$, initially placed outside the house: $out \neq locj$, for $j \in \{1,2,3,4\}$. As the phone can only access the house from its entrance, and $\delta = 1$, we have $\mathrm{d}(out, loc1) = 1$ and $\mathrm{d}(l, loci) \geq i$, for any $l \notin \{loc1, loc2, loc3, loc4\}$ and $i \in \{1,2,3,4\}$. Its interface $\mathtt{I}_P$ contains only one sensor, called *mode*, representing the touchscreen to control the boiler. This is a node-dependent sensor. The process *BoilerCtrl* reads sensor *mode* and forwards its value to the boiler manager in the patio, *BoilerMng*, via the Internet channel b ($\mathrm{rng}(b) = \infty$). The domain of the sensor *mode* is the set $\{\mathsf{man}, \mathsf{auto}\}$, where man stands for manual and auto for automatic; initially, $\mathtt{I}_P(mode) = \mathsf{auto}$. In *Phone* there is a second process, called *LightCtrl*, which allows the smartphone to switch on lights *only when* getting in touch with the light managers installed in the rooms. Here channels c_1 and c_2 serve to control the lights of entrance and lounge, respectively; these are short-range channels: $\mathrm{rng}(c_1) = \mathrm{rng}(c_2) = 0$.

The smart home *Home* consists of three stationary nodes: LM_1, LM_2 and BM.

The light managers processes $LightMng_1$ and $LightMng_2$, are placed in LM_1 and LM_2, respectively. They manage the corresponding lights via the actuators $light_j$, for $j \in \{1,2\}$. The domain of these actuators is $\{\mathsf{on}, \mathsf{off}\}$; initially, $\mathtt{I}_j(light_j) = \mathsf{off}$, for $j \in \{1,2\}$.

The boiler manager process *BoilerMng* is placed in BM (node n_B). Here, the physical interface $\mathtt{I}_B$ contains a sensor named *temp* and an actuator called *boiler*; *temp* is a location-dependent temperature sensor, whose domain is $\mathbb{N}$, and *boiler* is an actuator to display boiler functionality, whose domain is $\{\mathsf{on}, \mathsf{off}\}$. The boiler manager can work either in automatic or in manual mode. In automatic mode, sensor *temp* is periodically checked: if the temperature is under a threshold Θ then the boiler will be switched on, otherwise it will be switched off. Conversely, in manual mode, the boiler is always switched on. Initially, the boiler is in automatic mode, $\mathtt{I}_B(temp) = \Theta$, and $\mathtt{I}_B(boiler) = \mathsf{off}$.

Our system *Sys* enjoys a number of desirable *run-time properties*. For instance, if the boiler is in manual mode or its temperature is under the threshold Θ then the boiler will get switched on, within one time unit. Conversely, if the boiler is in automatic mode and its temperature is higher than or equal to the threshold Θ, then the boiler will get switched off within one time unit. In general, similar properties cannot be expressed in untimed calculi. Finally, our last property states the phone cannot act on the lights of the two rooms at the same time, manifesting a kind of "ubiquity".

For the sake of simplicity, in the following proposition we omit location names both in barbs and in sensor updates, writing $\downarrow_{a!v}$ instead of $\downarrow_{a@h!v}$, and $[s \mapsto v]$ instead of $[s@h \mapsto v]$. Furthermore, the system Sys' will denote an arbitrary (stable) derivative of *Sys*.

Proposition 3.1. *Let* $Sys\ (\rightarrow_{\mathrm{i}}^{*} \rightarrow_{\sigma})^{*}\ Sys'$, *for some* Sys'.

- *If* $Sys'[mode \mapsto \mathsf{man}] \rightarrow_{\mathrm{i}}^{*} Sys'' \rightarrow_{\sigma}$ *then* $Sys'' \downarrow_{boiler!\mathsf{on}}$.
- *If* $Sys'[temp \mapsto t] \rightarrow_{\mathrm{i}}^{*} Sys'' \rightarrow_{\sigma}$, *with* $t < \Theta$, *then* $Sys'' \downarrow_{boiler!\mathsf{on}}$.
- *If* $Sys'[temp \mapsto t] \rightarrow_{\mathrm{i}}^{*} Sys'' \rightarrow_{\sigma}$, *with* $t \geq \Theta$, *then* $Sys'' \downarrow_{boiler!\mathsf{off}}$.
- *If* $Sys' \rightarrow_{\mathrm{i}}^{*} Sys'' \downarrow_{light_1!\mathsf{on}}$ *then* $Sys'' \downarrow_{light_2!\mathsf{off}}$, *and vice versa.*

Table 5
Smart home in CaIT with a position-based light management.

$$\begin{array}{rcl}
\overline{Sys} & \stackrel{\text{def}}{=} & \overline{Phone} \mid \overline{Home} \\
\overline{Phone} & \stackrel{\text{def}}{=} & n_P[\mathtt{I}_P \bowtie BoilerCtrl \mid \overline{LightCtrl}]^{\mathrm{m}}_{out} \\
\overline{Home} & \stackrel{\text{def}}{=} & Home \mid \overline{CLM} \\
\overline{CLM} & \stackrel{\text{def}}{=} & n_C[\emptyset \bowtie \overline{CLightMng}]^{\mathrm{in}}_{loc3} \\
\overline{LightCtrl} & \stackrel{\text{def}}{=} & \mathsf{fix}\, X.@(x).\lfloor \overline{g}\langle x\rangle.\sigma.X \rfloor X \\
\overline{CLightMng} & \stackrel{\text{def}}{=} & \mathsf{fix}\, X.\lfloor g(y).[y = loc1]\lfloor \overline{c_1}\langle \mathsf{on}\rangle.\sigma.X \rfloor X; [y = loc4]\lfloor \overline{c_2}\langle \mathsf{on}\rangle.\sigma.X \rfloor X; \sigma.X \rfloor X
\end{array}$$

The proof of Proposition 3.1 can be found in the Appendix.

Finally, we propose a variant of our system, where lights functionality depends on the position of the smartphone. Intuitively, the smartphone detects its current GPS position, via the process $@(x).P$, and then sends it to a centralised light manager process of the house, via an Internet channel g. We implement that by adding a module $\overline{CLM}$ inside the patio, at location $loc3$, running a process $\overline{CLightMng}$. This process will interact with the local light managers $LightMng_1$ and $LightMng_2$ to switch on/off lights, depending on the GPS position received from the smartphone. The communication between centralised and local light managers will use the two short-range channels c_1 and c_2, with slightly different transmission ranges: $\mathrm{rng}(c_1) = 2$ and $\mathrm{rng}(c_2) = 1$. In Table 5, we provide the formalisation of this new variant where new components (with respect to those of Table 4) have been overlined.

Proposition 3.1 holds for system $\overline{Sys}$ as well. Actually, the two systems are closely related.

Proposition 3.2. *For $\delta = 1$, $(\boldsymbol{\nu}\tilde{c})Sys \cong (\boldsymbol{\nu}\tilde{c})(\boldsymbol{\nu}g)\overline{Sys}$.*

The bisimulation proof-technique developed in the remainder of the paper will be very useful to prove such kind of non-trivial system equalities (see Proposition 5.13 and Theorem 5.8).

We end this section with a comment. While reading this case study the reader should have realised that our reduction semantics does not model sensor updates. This is because sensor changes depend on the physical environment, while a reduction semantics models the evolution of a system in isolation. Interactions with the environment will be treated in the *extensional semantics* defined in the next section.

4. Labelled transition semantics

In this section, we provide two labelled semantic models, in the SOS style of Plotkin [12]: the *intensional semantics* and the *extensional semantics*. The adjective intensional is used to stress the fact that the actions of that semantics correspond to those activities which can be performed by a system in isolation, without any interaction with the external environment. Whereas, the extensional semantics focuses on those activities which require a contribution of the environment.

4.1. Intensional semantics

Since our syntax distinguishes between networks and processes, we have two different kinds of transitions:

- $P \xrightarrow{\lambda} Q$, with $\lambda \in \{\sigma, \tau, \overline{c}v, cv, @h, s?v, a!v\}$, for *process transitions*;
- $M \xrightarrow{\nu} N$, with $\nu \in \{\sigma, \tau, a, \overline{c}v@h, cv@h\}$, for *network transitions*.

In Table 6 we report standard transition rules for processes, very much in the style of [13]. Rules (SndP), (RcvP) and (Com) model communications along channel c. Rule (PosP) is for extracting the physical position of the embedding node. Rules (Sensor) and (Actuator) serve to read sensors, and to write on actuators, respectively. Rules (ParP) and (Fix) are straightforward. The remaining rules allow us the derive the timed action σ. In Rule (Delay) a timed prefix is consumed. Rule (Timeout) models timeouts when channel communication is not possible in the current time interval. Rule (TimeParP) is for time synchronisation of parallel processes. The symmetric counterparts of Rules (ParP) and (Com) are omitted. We recall that we assume $[b]P; Q = P$ if $[\![b]\!] = \mathsf{true}$, and $[b]P; Q = Q$ if $[\![b]\!] = \mathsf{false}$.

In Table 7 we report the transition rules for networks. Rule (Pos) extracts the position of a node. Rule (SensRead) models the reading of a value from a sensor of the enclosing node. Rules (ActUnChg) and (ActChg) describe the writing of a value v on an actuator a of the node, distinguishing whether the value of the actuator is changed or not. Rule (LocCom) models intra-node communications. Rule (TimeStat) models the passage of time for a stationary node. Rule (TimeMob) models both time passing and node mobility at the end of a time interval. Rules (Snd) and (Rcv) represent transmission and reception along a global channel. Rule (GlbCom) models inter-node communications. The remaining rules are straightforward. The symmetric counterparts of Rules (ParN) and (GlobCom) are omitted.

As expected, the reduction semantics and the labelled intensional semantics coincide.

Please cite this article in press as: R. Lanotte, M. Merro, A semantic theory of the Internet of Things, Inf. Comput. (2018), https://doi.org/10.1016/j.ic.2018.01.001

Table 6
Intensional semantics for processes.

$$\text{(SndP)}\ \frac{-}{\lfloor \overline{c}\langle v\rangle.P \rfloor Q \xrightarrow{\overline{c}v} P} \qquad \text{(RcvP)}\ \frac{-}{\lfloor c(x).P \rfloor Q \xrightarrow{cv} P\{^v/_x\}}$$

$$\text{(PosP)}\ \frac{-}{@(x).P \xrightarrow{@h} P\{^h/_x\}} \qquad \text{(Com)}\ \frac{P \xrightarrow{\overline{c}v} P' \quad Q \xrightarrow{cv} Q' \quad \mathrm{rng}(c) = -1}{P \mid Q \xrightarrow{\tau} P' \mid Q'}$$

$$\text{(Sensor)}\ \frac{-}{s?(x).P \xrightarrow{s?v} P\{^v/_x\}} \qquad \text{(Actuator)}\ \frac{-}{a!v.P \xrightarrow{a!v} P}$$

$$\text{(ParP)}\ \frac{P \xrightarrow{\lambda} P' \quad \lambda \neq \sigma}{P \mid Q \xrightarrow{\lambda} P' \mid Q} \qquad \text{(Fix)}\ \frac{P\{^{\mathsf{fix}\,X.P}/_X\} \xrightarrow{\lambda} Q}{\mathsf{fix}\,X.P \xrightarrow{\lambda} Q}$$

$$\text{(TimeNil)}\ \frac{-}{\mathsf{nil} \xrightarrow{\sigma} \mathsf{nil}} \qquad \text{(Delay)}\ \frac{-}{\sigma.P \xrightarrow{\sigma} P}$$

$$\text{(Timeout)}\ \frac{-}{\lfloor \pi.P \rfloor Q \xrightarrow{\sigma} Q} \qquad \text{(TimeParP)}\ \frac{P \xrightarrow{\sigma} P' \quad Q \xrightarrow{\sigma} Q' \quad P|Q \xrightarrow{\tau}\!\!\!\!\!/}{P \mid Q \xrightarrow{\sigma} P' \mid Q'}$$

Table 7
Intensional semantics for networks.

$$\text{(Pos)}\ \frac{P \xrightarrow{@h} P'}{n[\mathtt{I} \bowtie P]^{\mu}_{h} \xrightarrow{\tau} n[\mathtt{I} \bowtie P']^{\mu}_{h}} \qquad \text{(SensRead)}\ \frac{\mathtt{I}(s) = v \quad P \xrightarrow{s?v} P'}{n[\mathtt{I} \bowtie P]^{\mu}_{h} \xrightarrow{\tau} n[\mathtt{I} \bowtie P']^{\mu}_{h}}$$

$$\text{(ActUnChg)}\ \frac{\mathtt{I}(a) = v \quad P \xrightarrow{a!v} P'}{n[\mathtt{I} \bowtie P]^{\mu}_{h} \xrightarrow{\tau} n[\mathtt{I} \bowtie P']^{\mu}_{h}} \qquad \text{(LocCom)}\ \frac{P \xrightarrow{\tau} P'}{n[\mathtt{I} \bowtie P]^{\mu}_{h} \xrightarrow{\tau} n[\mathtt{I} \bowtie P']^{\mu}_{h}}$$

$$\text{(ActChg)}\ \frac{\mathtt{I}(a) \neq v \quad P \xrightarrow{a!v} P' \quad \mathtt{I}' := \mathtt{I}[a \mapsto v]}{n[\mathtt{I} \bowtie P]^{\mu}_{h} \xrightarrow{a} n[\mathtt{I}' \bowtie P']^{\mu}_{h}}$$

$$\text{(TimeStat)}\ \frac{P \xrightarrow{\sigma} P' \quad n[\mathtt{I} \bowtie P]^{\mathrm{s}}_{h} \xrightarrow{\tau}\!\!\!\!\!/}{n[\mathtt{I} \bowtie P]^{\mathrm{s}}_{h} \xrightarrow{\sigma} n[\mathtt{I} \bowtie P']^{\mathrm{s}}_{h}} \qquad \text{(TimeMob)}\ \frac{P \xrightarrow{\sigma} P' \quad n[\mathtt{I} \bowtie P]^{\mathrm{m}}_{h} \xrightarrow{\tau}\!\!\!\!\!/ \quad \mathrm{d}(h,k) \leq \delta}{n[\mathtt{I} \bowtie P]^{\mathrm{m}}_{h} \xrightarrow{\sigma} n[\mathtt{I} \bowtie P']^{\mathrm{m}}_{k}}$$

$$\text{(Snd)}\ \frac{P \xrightarrow{\overline{c}v} P' \quad \mathrm{rng}(c) \geq 0}{n[\mathtt{I} \bowtie P]^{\mu}_{h} \xrightarrow{\overline{c}v@h} n[\mathtt{I} \bowtie P']^{\mu}_{h}} \qquad \text{(Rcv)}\ \frac{P \xrightarrow{cv} P' \quad \mathrm{rng}(c) \geq 0}{n[\mathtt{I} \bowtie P]^{\mu}_{h} \xrightarrow{cv@h} n[\mathtt{I} \bowtie P']^{\mu}_{h}}$$

$$\text{(GlbCom)}\ \frac{M \xrightarrow{\overline{c}v@k} M' \quad N \xrightarrow{cv@h} N' \quad \mathrm{d}(h,k) \leq \mathrm{rng}(c)}{M|N \xrightarrow{\tau} M'|N'}$$

$$\text{(ParN)}\ \frac{M \xrightarrow{\nu} M' \quad \nu \neq \sigma}{M|N \xrightarrow{\nu} M'|N} \qquad \text{(TimePar)}\ \frac{M \xrightarrow{\sigma} M' \quad N \xrightarrow{\sigma} N' \quad M|N \xrightarrow{\tau}\!\!\!\!\!/}{M|N \xrightarrow{\sigma} M'|N'}$$

$$\text{(TimeZero)}\ \frac{-}{\mathbf{0} \xrightarrow{\sigma} \mathbf{0}} \qquad \text{(Res)}\ \frac{M \xrightarrow{\nu} N \quad \nu \notin \{\overline{c}v@h, cv@h\}}{(\boldsymbol{\nu} c)M \xrightarrow{\nu} (\boldsymbol{\nu} c)N}$$

Theorem 4.1 *(Harmony theorem). Let $\omega \in \{\tau, a, \sigma\}$:*

- *$M \xrightarrow{\omega} M'$ implies $M \rightarrow\!\!\triangleright_{\omega} M'$;*
- *$M \rightarrow\!\!\triangleright_{\omega} M'$ implies $M \xrightarrow{\omega} M''$ for some M'' such that $M' \equiv M''$.*

4.2. Extensional semantics

Here we redesign our LTS to focus on the interactions of our (networks of) systems with the external environment. As the environment has a *logical part* (the parallel nodes) and a *physical part* (the physical world) our extensional semantics distinguishes two different kinds of transitions:

- $M \xrightarrow{\alpha} N$, *logical transitions*, for $\alpha \in \{\tau, \sigma, a, \overline{c}v \triangleright k, cv \triangleright k\}$, to denote the interaction with the *logical environment*; here, actuator changes, τ- and σ-actions are inherited from the intensional semantics, so we don't provide inference rules for them;

Please cite this article in press as: R. Lanotte, M. Merro, A semantic theory of the Internet of Things, Inf. Comput. (2018), https://doi.org/10.1016/j.ic.2018.01.001

Table 8
Extensional semantics: additional rules.

$$(\text{SndObs})\ \frac{M \xrightarrow{\overline{c}v@h} M' \quad \mathrm{d}(h,k) \leq \mathrm{rng}(c)}{M \xrightarrow{\overline{c}v \triangleright k} M'} \qquad (\text{RcvObs})\ \frac{M \xrightarrow{cv@h} M' \quad \mathrm{d}(k,h) \leq \mathrm{rng}(c)}{M \xrightarrow{cv \triangleright k} M'}$$

$$(\text{SensEnv})\ \frac{v \text{ in the domain of } s}{M \xrightarrow{s@h?v} M[s@h \mapsto v]} \qquad (\text{ActEnv})\ \frac{M \downarrow_{a@h!v}}{M \xrightarrow{a@h!v} M}$$

- $M \xrightarrow{\alpha} N$, *physical transitions*, for $\alpha \in \{s@h?v, a@h!v\}$, to denote the interaction with the *physical world*, via sensors and actuators.

In Table 8 the extensional actions deriving from rules (SndObs) and (RcvObs) mention the location k of the logical environment which can *observe* the communication occurring at channel c. Rules (SensEnv) and (ActEnv) model the interaction of a network M with the physical environment. In particular, the environment can *nondeterministically update* the current value of a (location-dependent or node-dependent) sensor s with a value v, and can read the value v appearing on an actuator a at h. As already discussed in Section 2.2 the environment is potentially aware of its position when performing these actions.

Note that our LTSs are *image finite*. They are also *finitely branching*, and hence potentially *mechanisable*, under the obvious assumption of finiteness of all domains of admissible values, and the set of physical locations.

5. Coinductive characterisation

Based on our extensional semantics, we are ready to define a notion of weak bisimilarity which will be showed to be both sound and complete with respect to our contextual equivalence.

We adopt a standard notation for weak transitions. We denote with $\Longrightarrow$ the reflexive and transitive closure of τ-actions, namely $(\xrightarrow{\tau})^*$, whereas $\stackrel{\alpha}{\Longrightarrow}$ means $\Longrightarrow\xrightarrow{\alpha}\Longrightarrow$, and finally $\stackrel{\hat{\alpha}}{\Longrightarrow}$ denotes $\Longrightarrow$ if $\alpha = \tau$ and $\stackrel{\alpha}{\Longrightarrow}$ otherwise.

Definition 5.1 *(Bisimulation).* A binary symmetric relation $\mathcal{R}$ over networks is a *bisimulation* if $M\,\mathcal{R}\,N$ and $M \xrightarrow{\alpha} M'$ imply there exists N' such that $N \stackrel{\hat{\alpha}}{\Longrightarrow} N'$ and $M'\,\mathcal{R}\,N'$. We say that M and N are *bisimilar*, written $M \approx N$, if $M\,\mathcal{R}\,N$ for some bisimulation $\mathcal{R}$.

A crucial result is that our bisimilarity is a congruence. In order to prove that, we need a technical lemma saying that the operator for sensor updates does not affect the number of consecutive instantaneous reductions of a network.

Definition 5.2. Let redi() be a function that given a network M returns an upper bound to the number of consecutive instantaneous reductions that M may perform:

$$\begin{array}{rcl@{\qquad}rcl} \mathrm{redi}(\mathbf{0}) & = & 0 & \mathrm{redi}(n[\mathtt{I} \bowtie P]^{\mu}_{h}) & = & \mathrm{pfxi}(P) \\ \mathrm{redi}((\boldsymbol{\nu} c)M) & = & \mathrm{redi}(M) & \mathrm{redi}(M|N) & = & \mathrm{redi}(M) + \mathrm{redi}(N) \end{array}$$

where pfxi() is a function that given a process P returns an upper bound to the number of untimed prefixes in P that give rise to an instantaneous reduction when P is plugged into a node (for details, see Definition A.4 in the Appendix).

Lemma 5.3. *For any network M, sensor s, location h and value v in the domain of s, it follows that* $\mathrm{redi}(M) = \mathrm{redi}(M[s@h \mapsto v])$.

Theorem 5.4. *The relation $\approx$ is contextual.*

Proof. We have to show that $\approx$ is preserved by parallel composition, channel restriction, and sensor updates. The most difficult part is to prove that $M \approx N$ entails $M[s@h \mapsto v] \approx N[s@h \mapsto v]$, for all sensors s, locations h, and values v in the domain of s. In fact, a standard approach to this proof, consisting in trying to show that the relation $\{\big(M[s@h \mapsto v]\,,\, N[s@h \mapsto v]\big) : M \approx N\}$ is a bisimulation, is not affordable. In fact, in general, if $M \xrightarrow{\tau} M'$ then we do not necessarily have $M[s@h \mapsto v] \xrightarrow{\tau} M'[s@h \mapsto v]$. This is because sensor updates may have an influence on the evolution of M.

Let $\succ$ be the well-founded relation over pairs of networks such that $(M, N) \succ (M', N')$ if and only if (i) $M \approx N$; (ii) $M' \approx N'$; (iii) $\mathrm{redi}(M) + \mathrm{redi}(N) > \mathrm{redi}(M') + \mathrm{redi}(N')$. Note that $\succ$ is trivially irreflexive. Moreover, redi(M) always returns a finite and positive integer, for any M (see Lemma A.6 in the Appendix). Thus, the relation $\succ$ does not have infinite descending chains and it is a well-founded relation.

The proof is by *well-founded induction* on the definition of $\succ$.

Base case. Let M and N be such that $M \approx N$ and $\mathrm{redi}(M) + \mathrm{redi}(N) = 0$. By Definition 5.2 and by inspection of the reduction semantics in Table 3, $\mathrm{redi}(N) = 0$ entails $N \not\rightarrow_{\mathrm{i}}$. In particular, we have $N \not\rightarrow_{\tau}$. By Theorem 4.1 it follows that

$N \stackrel{\tau}{\not\rightarrow}$. By an application of rule (SensEnv) we derive $M \xrightarrow{s@h?v} M[s@h \mapsto v]$. As $M \approx N$ there are N_1, N_2 and N' such that $N \Rightarrow N_1 \xrightarrow{s@h?v} N_1[s@h \mapsto v] = N_2 \Rightarrow N'$ with $M[s@h \mapsto v] \approx N'$. However, as $N \stackrel{\tau}{\not\rightarrow}$ it follows that $N = N_1$. By Lemma 5.3, $N_1 \stackrel{\tau}{\not\rightarrow}$ entails $N_1[s@h \mapsto v] \stackrel{\tau}{\not\rightarrow}$. This is enough to derive that $N' = N[s@h \mapsto v]$. And hence $M[s@h \mapsto v] \approx N[s@h \mapsto v]$.

Inductive Case. Let $M \approx N$. Without loss of generality we can assume $\mathrm{redi}(M) \geq \mathrm{redi}(N)$. Let $M \xrightarrow{s@h?v} M[s@h \mapsto v]$. As $M \approx N$ there is N' such that $N \stackrel{s@h?v}{\Longrightarrow} N'$ and $M[s@h \mapsto v] \approx N'$.

Now, if the number of τ-actions contained in the weak transition $N \stackrel{s@h?v}{\Longrightarrow} N'$ is 0, then $N \xrightarrow{s@h?v} N[s@h \mapsto v]$, with $M[s@h \mapsto v] \approx N[s@h \mapsto v]$, and there is nothing else to prove. Otherwise, we have $\mathrm{redi}(N') < \mathrm{redi}(N)$. By Lemma 5.3 it follows that $\mathrm{redi}(M) = \mathrm{redi}(M[s@h \mapsto v])$. Thus, $\mathrm{redi}(M[s@h \mapsto v]) + \mathrm{redi}(N') < \mathrm{redi}(M) + \mathrm{redi}(N)$. Hence $(M[s@h \mapsto v], N') \prec (M, N)$. By inductive hypothesis we know that $M[s@h \mapsto v][r@k \mapsto w] \approx N'[r@k \mapsto w]$ for any sensor r, location k and value w in the domain of r. Thus, if we choose $r = s$, $k = h$ and w the value such that $M[s@h \mapsto v][s@h \mapsto w] = M[s@h \mapsto w] = M$, then we get $M \approx N'[s@h \mapsto w]$.[3]

Finally, we need two small sub-results to conclude the proof.

- By Lemma 5.3 we have $\mathrm{redi}(N'[s@h \mapsto w]) = \mathrm{redi}(N')$. Since $\mathrm{redi}(N') < \mathrm{redi}(N)$ it follows that $(M, N'[s@h \mapsto w]) \prec (M, N)$. By inductive hypothesis we can close under the operator $[s@h \mapsto v]$, getting $M[s@h \mapsto v] \approx N'[s@h \mapsto w][s@h \mapsto v]$.
- Since $\approx$ is a transitive relation, $M \approx N'[s@h \mapsto w]$ and $M \approx N$, we derive that $N \approx N'[s@h \mapsto w]$. Since, $\mathrm{redi}(N) \leq \mathrm{redi}(M)$ (this was an initial assumption) and $\mathrm{redi}(N'[s@h \mapsto w]) = \mathrm{redi}(N') < \mathrm{redi}(N)$ it follows that $(N, N'[s@h \mapsto w]) \prec (M, N)$. By inductive hypothesis we can derive $N[s@h \mapsto v] \approx N'[s@h \mapsto w][s@h \mapsto v]$.

From these two facts, and by transitivity of $\approx$ we finally get $M[s@h \mapsto v] \approx N[s@h \mapsto v]$.

The proof that $\approx$ is preserved by parallel composition and channel restriction can be found in the Appendix. □

In order to prove that our labelled bisimilarity is sound with respect to reduction barbed congruence, we need the following technical result relating barbs and extensional actions.

Proposition 5.5. *$M \downarrow_{a@h!v}$ if and only if $M \xrightarrow{a@h!v} M$.*

Proof. It follows from the definition of rule (ActEnv). □

Theorem 5.6 *(Soundness). Let M and N be two networks such that $M \approx N$, then $M \cong N$.*

Proof. We recall that $\cong$ is defined as the largest symmetric reduction which is reduction closed, barb preserving and contextual.

Let us prove that the labelled bisimilarity is reduction closed. Suppose that $M \rightarrowtriangle M'$. Then we have two cases: either $M \rightarrowtriangle_\tau M'$ or $M \rightarrowtriangle_\sigma M'$. In the first case Theorem 4.1 implies $M \xrightarrow{\tau} \equiv M'$. As $M \approx N$ there exists N' such that $N \Rightarrow N'$ and $M' \approx N'$. Now, by Theorem 4.1 we have that each of the τ-actions in the sequence $N \Rightarrow N'$ can be rewritten in terms of $\rightarrowtriangle_\tau$. Thus the entire sequence $N \Rightarrow N'$ can be rewritten as a sequence of instantaneous reductions $N \rightarrowtriangle_\tau^* N'$, which is a particular case of $N \rightarrowtriangle^* N'$, with $N' \approx M'$. Let us consider now the second case: $M \rightarrowtriangle_\sigma M'$. By Theorem 4.1 it follows that $M \xrightarrow{\sigma} \equiv M'$. As $M \approx N$ there exists N' such that $N \stackrel{\sigma}{\Longrightarrow} N'$ and $N' \approx M'$. By several applications of Theorem 4.1 we get $N \rightarrowtriangle_\tau^* \rightarrowtriangle_\sigma \rightarrowtriangle_\tau^* N'$. Thus, $N \rightarrowtriangle^* N'$, with $M' \approx N'$.

The proof that $\approx$ is closed with respect to $M \rightarrowtriangle_a N$ is similar.

From reduction closure and Proposition 5.5 it follows immediately that $\approx$ is barb preserving.

Theorem 5.4 says that our labelled bisimilarity is contextual.

As $\cong$ is the largest relation which is reduction closed, barb-preserving and contextual, it follows that $\approx \subseteq \cong$. □

Next, we prove now that our bisimilarity is also complete. The proof relies on showing that for each extensional action α it is possible to exhibit a test T_α which determines whether or not a system M can perform the action α. We need a technical lemma to cut down observing contexts.

Lemma 5.7. *Let M and N be two networks. Let $O = n[\mathtt{I} \bowtie a!v.\mathrm{nil}]_k^s$, for an arbitrary node name n, an arbitrary actuator a, and arbitrary values v and w, in the domain of a, such that $\mathtt{I}$ is only defined for a and $\mathtt{I}(a) = w \neq v$. If both $M|O$ and $N|O$ are well-formed and $M|O \cong N|O$ then $M \cong N$.*

Theorem 5.8 *(Completeness). Let M and N such that $M \cong N$, then $M \approx N$.*

[3] By Definition 2.9 the value on M of the sensor s located at h must be w, if defined. Otherwise it can be any admissible value for s.

Proof. We show that relation $\mathcal{R} = \{(M, N) \mid M \cong N\}$ is a bisimulation up to $\equiv$.[4] Let us consider two networks M and N such that $(M, N) \in \mathcal{R}$. We proceed by case analysis on the possible extensional actions of M (the case when N moves first is similar).

First, we consider *logical transitions*.

- Let us suppose that $M \xrightarrow{a} M'$. By Theorem 4.1 we derive $M \rightarrowtriangle_a M'$. Let us define the test T_a:

  $$T_a \stackrel{\text{def}}{=} n[\jmath \bowtie b!1]^{\mathrm{s}}_k$$

  where n is a fresh node name and b is a fresh actuator such that $\jmath(b) = 0$. By Proposition 2.3, no σ-move can fire if a reduction $\rightarrowtriangle_b$ is possible. Thus, the presence of a barb $\Downarrow_{b@k!0}$ means that no σ-actions have occurred yet. Since $M \rightarrowtriangle_a M'$, we can apply rule (parn) to infer $M|T_a \rightarrowtriangle_a M'|T_a$, with $M'|T_a \downarrow_{b@k!0}$. As $M \cong N$ and the relation $\cong$ is both contextual and reduction closed, it follows that $N|T_a \rightarrowtriangle^* \rightarrowtriangle_a \rightarrowtriangle^* \hat{N}$, for some $\hat{N}$, with $M'|T_a \cong \hat{N}$. As a consequence, $\hat{N} \Downarrow_{b@k!0}$. This implies that $\hat{N} \equiv N'|T_a$ for some N', such that $N|T_a \rightarrowtriangle^* \rightarrowtriangle_a \rightarrowtriangle^* N'|T_a$, with $N \rightarrowtriangle^* \rightarrowtriangle_a \rightarrowtriangle^* N'$, and $M'|T_a \cong N'|T_a$. As the presence of a barb $\Downarrow_{b@k!0}$ ensures that no σ-actions have occurred, it follows that $N \rightarrowtriangle^*_\tau \rightarrowtriangle_a \rightarrowtriangle^*_\tau N'$. By several applications of Theorem 4.1 it follows that $N \stackrel{a}{\Longrightarrow} \equiv N'$ (this relies on the straightforward result that $\equiv$ is a strong bisimulation). By $M'|T_a \cong N'|T_a$ and Lemma 5.7 we derive $M' \cong N'$. This implies that $M' \equiv \mathcal{R} \equiv N'$.
- Let us suppose that $M \xrightarrow{\tau} M'$. This case is similar to the previous one with $T_\tau = T_a$.
- Let us suppose that $M \xrightarrow{\sigma} M'$. By Theorem 4.1 we derive $M \rightarrowtriangle_\sigma M'$. As $M \cong N$, by Proposition 2.14 there exists N' such that $N \rightarrowtriangle^*_\tau \rightarrowtriangle_\sigma \rightarrowtriangle^*_\tau N'$ and $M' \cong N'$. By several applications of Theorem 4.1 we obtain $N \stackrel{\sigma}{\Longrightarrow} \equiv N'$. As $M' \cong N'$, it follows that $M' \equiv \mathcal{R} \equiv N'$.
- Let us suppose that $M \xrightarrow{\overline{c}v \triangleright k} M'$. This transition can only be derived by an application of rule (SndObs) if $M \xrightarrow{\overline{c}v@h} M'$, for some h, such that $\mathrm{d}(h, k) \leq \mathrm{rng}(c)$. Let us build up a context that is capable to observe the action $\overline{c}v \triangleright k$. We define testing term $T_{\overline{c}v \triangleright k}$. For simplicity, in the following we abbreviate it with T:

  $$T \stackrel{\text{def}}{=} m[\jmath \bowtie \lfloor c(x).[x = v]b!1.b!0; \mathsf{nil} \rfloor \mathsf{nil}]^{\mathrm{s}}_k$$

  where m is a fresh node name and b is a fresh actuator name such that $\jmath(b) = 0$. The intuition behind this testing process is the following: T has barb $\Downarrow_{b@k!1}$ only if the communication along c has already occurred and no time actions have been fired (Proposition 2.3).

  Since $\cong$ is contextual, $M \cong N$ implies $M|T \cong N|T$. From $M \xrightarrow{\overline{c}v@h} M'$ we can easily infer $M|T \xrightarrow{\tau} \xrightarrow{b} M'|T'$, with $T' = m[\jmath[b \mapsto 1] \bowtie b!0.\mathsf{nil}]^{\mathrm{s}}_k$. Notice that $M'|T' \downarrow_{b@k!1}$. By Theorem 4.1, we derive $M|T \rightarrowtriangle_\tau \rightarrowtriangle_b M'|T'$. As $M|T \cong N|T$ it follows that $N|T \rightarrowtriangle^* \rightarrowtriangle_b \rightarrowtriangle^* \hat{N}$, with $\hat{N} \Downarrow_{b@k!1}$. This implies that $\hat{N} \equiv N'|T'$, for some N'. Furthermore, no timed actions have occurred in the reduction sequence, and hence: $N|T \rightarrowtriangle^*_\tau \rightarrowtriangle_b \rightarrowtriangle^*_\tau N'|T'$. By several applications of Theorem 4.1 we obtain $N|T \Rightarrow \xrightarrow{b} \Rightarrow \equiv N'|T'$. This implies that $N \stackrel{\overline{c}v@h'}{\Longrightarrow} \equiv N'$, for some h' such that $\mathrm{d}(h', k) \leq \mathrm{rng}(c)$. By an application of rule (SndObs) we get $N \stackrel{\overline{c}v \triangleright k}{\Longrightarrow} \equiv N'$. From $M'|T \cong N'|T$ and Lemma 5.7 we derive $M' \cong N'$. This allows us to derive that $M' \equiv \mathcal{R} \equiv N'$.
- The case $M \xrightarrow{cv \triangleright k} M'$ is similar to the previous one. The observing term is

  $$T_{cv \triangleright k} \stackrel{\text{def}}{=} m[\jmath \bowtie \lfloor \overline{c}\langle v \rangle.b!1.b!0.\mathsf{nil} \rfloor \mathsf{nil}]^{\mathrm{s}}_k$$

  where m is a fresh node name and b is a fresh actuator name such that $\jmath(b) = 0$.

Let us consider now *physical transitions*. Here, as already explained in Section 2.3, we will not provide an observing context as our language for IoT systems does not allow us to write physical observers.

- Let $M \xrightarrow{a@h!v} M'$. Since this transition can be only derived by an application of rule (ActEnv), it follows that $M' = M$ and $M \downarrow_{a@h!v}$. By Proposition 2.16 we obtain $N \downarrow_{a@h!v}$. By applying again rule (ActRead) to N, we obtain $N \xrightarrow{a@h!v} N' = N$ with $(M', N') \in \mathcal{R}$.
- Let $M \xrightarrow{s@h?v} M'$. Since this transition can be only derived by an application of rule (SenEnv), it follows that $M' = M[s@h \mapsto v]$. By an application of the same rule (SensEnv) we obtain $N \xrightarrow{s@h?v} N' = N[s@h \mapsto v]$. As $\cong$ is contextual we have $M[s@h \mapsto v] \cong N[s@h \mapsto v]$. This implies that $(M', N') \in \mathcal{R}$. □

By Theorem 5.6 and Theorem 5.8 we derive the full abstraction result: reduction barbed congruence coincides with our labelled bisimilarity.

[4] We recall that two structural congruent networks have exactly the same labelled transitions [10].

Please cite this article in press as: R. Lanotte, M. Merro, A semantic theory of the Internet of Things, Inf. Comput. (2018), https://doi.org/10.1016/j.ic.2018.01.001

Theorem 5.9 *(Full abstraction). $M \approx N$ if and only if $M \cong N$.*

Remark 5.10. A consequence of Theorem 5.9 and Remark 2.12 is that our bisimulation proof-technique remains sound in a setting where nondeterministic sensor updates are replaced by some specific model for sensors.

5.1. Algebraic laws and examples

As testbed for our bisimulation proof-technique we prove a number of algebraic laws on well-formed networks. Some of these laws are valid with respect to a stronger form of bisimilarity which takes into account the number of τ-actions performed by a process. The *expansion* relation [23], written $\lesssim$, is an asymmetric variant of $\approx$ such that $P \lesssim Q$ holds if $P \approx Q$ and Q has at least as many τ-moves as P.

Theorem 5.11 *(Some algebraic laws).*

1. $n[\mathtt{I} \bowtie a!v.P|R]_h^\mu \gtrsim n[\mathtt{I} \bowtie P|R]_h^\mu$, *if* $\mathtt{I}(a) = v$ *and* a *does not occur in* R;
2. $n[\mathtt{I} \bowtie @(x).P|R]_h^\mu \gtrsim n[\mathtt{I} \bowtie \{^h/_x\}P|R]_h^\mu$;
3. $n[\mathtt{I} \bowtie \lfloor\overline{c}\langle v\rangle.P\rfloor S|\lfloor c(x).Q\rfloor T|R]_h^\mu \gtrsim n[\mathtt{I} \bowtie P|Q\{^v/_x\}|R]_h^\mu$, *if* c *is not in* R *and* $\mathrm{rng}(c) = -1$;
4. $(\nu c)(n[\mathtt{I} \bowtie \lfloor\overline{c}\langle v\rangle.P\rfloor S|R]_h^\mu \mid m[\mathtt{J} \bowtie \lfloor c(x).Q\rfloor T|U]_k^{\mu'}) \gtrsim (\nu c)(n[\mathtt{I} \bowtie P|R]_h^\mu \mid m[\mathtt{J} \bowtie Q\{^v/_x\}|U]_k^{\mu'})$, *if* $\mathrm{rng}(c) = \infty$ *and* c *does not occur in* R *and* U;
5. $n[\mathtt{I} \bowtie P]_h^\mu \approx n[\mathtt{I} \bowtie \mathsf{nil}]_h^\mu$, *if subterms* $\lfloor\pi.P_1\rfloor P_2$ *or* $a!v.P_1$ *do not occur in* P;
6. $n[\mathtt{I} \bowtie \mathsf{nil}]_h^\mu \approx \mathbf{0}$, *if* $\mathtt{I}(a)$ *is undefined for any actuator* a;
7. $n[\emptyset \bowtie P]_h^{\mathrm{m}} \approx m[\emptyset \bowtie P]_k^{\mathrm{s}}$, *if* P *does not contain terms of the form* $@(x).Q$ *and for any channel* c *in* P *either* $\mathrm{rng}(c) = \infty$ *or* $\mathrm{rng}(c) = -1$.

Laws 1–4 are a sort of tau-laws. Laws 5 and 6 model garbage collection of processes and nodes, respectively. Law 7 gives a sufficient condition for node anonymity as well as for non-observable node mobility.

Next, we show that our labelled bisimilarity can be used to deal with more complicated systems. In the following, we apply non-trivial up to expansion proof-techniques[5] to formally prove that the two systems *Sys* and $\overline{Sys}$ mentioned in Proposition 3.2 are bisimilar (up to an obvious channel restriction). In this respect, the first four laws of Theorem 5.11 are fundamentals.

The following lemma basically says that, under specific conditions, our bisimilarity is preserved by parallel composition on processes within nodes.

Lemma 5.12. *If* $(\nu\tilde{c})\left(n[\mathtt{I} \bowtie P_1]_h^\mu | O_1\right) \approx (\nu\tilde{d})\left(n[\mathtt{I} \bowtie P_2]_k^\mu | O_2\right)$ *then* $(\nu\tilde{c})\left(n[\mathtt{I} \bowtie P_1|R]_h^\mu | O_1\right) \approx (\nu\tilde{d})\left(n[\mathtt{I} \bowtie P_2|R]_k^\mu | O_2\right)$, *for any process* R *that can only read sensors, transmit along some fresh Internet channel, and let time pass.*

We can now rely on Theorem 5.9 to rephrase Proposition 3.2 by replacing reduction barbed congruence with our labelled bisimilarity.

Proposition 5.13. *If* $\delta = 1$ *then* $(\nu\tilde{c})Sys \approx (\nu\tilde{c})(\nu g)\overline{Sys}$.

Proof. First of all, notice that we can focus on smaller systems. This is because:

- $(\nu\tilde{c})Sys = (\nu\tilde{c})\left(Phone|LM_1|BM|LM_2\right) \equiv (\nu\tilde{c})\left(Phone|LM_1|LM_2\right) \mid BM$
- $(\nu\tilde{c}, g)\overline{Sys} = (\nu\tilde{c}, g)\left(\overline{Phone}|LM_1|BM|LM_2|\overline{CLM}\right) \equiv (\nu\tilde{c}, g)\left(\overline{Phone}|LM_1|LM_2|\overline{CLM}\right) \mid BM$.

By Theorem 5.4 the relation $\approx$ is preserved by parallel composition. Thus, in order to prove our result it is enough to show that:

$$(\nu\tilde{c})\left(Phone|LM_1|LM_2\right) \approx (\nu\tilde{c})(\nu g)\left(\overline{Phone}|LM_1|LM_2|\overline{CLM}\right).$$

Actually, we can consider even smaller systems. As channels c_1, c_2 and g do not occur in *BoilerCtrl*, and $Phone = n_P[\mathtt{I}_P \bowtie BoilerCtrl \mid LightCtrl]_{out}^{\mathrm{m}}$, and $\overline{Phone} = n_P[\mathtt{I}_P \bowtie BoilerCtrl \mid \overline{LightCtrl}]_{out}^{\mathrm{m}}$, by Lemma 5.12 it is enough to show that:

$$(\nu\tilde{c})\left(n_P[\mathtt{I}_P \bowtie LightCtrl]_{out}^{\mathrm{m}}|LM_1|LM_2\right) \approx (\nu\tilde{c})(\nu g)\left(n_P[\mathtt{I}_P \bowtie \overline{LightCtrl}]_{out}^{\mathrm{m}}|LM_1|LM_2|\overline{CLM}\right).$$

Let us call S_{L} the system on the left side, and S_{R} the system on the right side. We define the relation:

[5] The up-to-expansion proof-technique for the standard notion of weak bisimilarity is notoriously sound [10].

Please cite this article in press as: R. Lanotte, M. Merro, A semantic theory of the Internet of Things, Inf. Comput. (2018), https://doi.org/10.1016/j.ic.2018.01.001

$$\mathcal{R} \stackrel{\text{def}}{=} \bigcup_{i=1}^{17} \Big((\nu \tilde{c}) M_i \,,\, (\nu \tilde{c})(\nu g) N_i \Big)$$

where the details of the pairs (M_i, N_i) are given in the Appendix. We will prove that the symmetric closure of $\mathcal{R}$ is a bisimulation up to expansion [10]. Then, we will show that $S_{\mathrm{L}} = (\nu \tilde{c}) M_1 \, \mathcal{R} \, (\nu \tilde{c})(\nu g) N_1 \lesssim S_{\mathrm{R}}$. As the up to expansion technique is sound, and the expansion relation $\lesssim$ is a transitive relation, it follows that $S_{\mathrm{L}} \approx S_{\mathrm{R}}$. The details of the proof can be found in the Appendix. □

6. Conclusions and related work

We have proposed a process calculus, called CaIT, to investigate the semantic theory of networked systems in the Internet of Things paradigm. The dynamics of CaIT is formalised by means of an intuitive reduction semantics and (a more operational) labelled intensional semantics that model the evolution of systems in isolation. An Harmony theorem shows that these two different operational semantics coincide. An extensional semantics has then been defined to emphasise the interaction of IoT systems with the environment. The extensional semantics has been used to define a labelled bisimilarity which has been proved to be a coinductive characterisation of a natural notion of contextual equivalence. Our bisimilarity has been used to prove non-trivial system equalities.

To our knowledge, Lanese et al.'s IoT-calculus [8] is the first process calculus for IoT systems capturing the interaction between sensors, actuators and computing processes. Smart objects are represented as point-to-point communicating nodes of heterogeneous networks. The network topology is represented as a graph whose links can be nondeterministically established or destroyed. The paper contains a labelled transition system with two different kinds of transitions. The first one takes into account interactions with the physical environment, similarly to our physical transitions, but includes also topology changes. The second kind of transition models nodes activities, mainly communications, similarly to our logical transitions. The paper proposes two notions of bisimilarity: one using only the first kind of transitions and equating systems from the point of view of the end user, and a second one using all transitions and equating systems from the point of view of the other devices.

We report here the main differences between CaIT and the IoT-calculus. In CaIT, we support timed behaviours, with desirable time, consistency and fairness properties (see, for instance, Proposition 3.1). Both sensors and actuators in CaIT are under the control of a single entity, i.e. the controller process of the node where they are deployed. This was a security issue. The nondeterministic link entailment of the IoT-calculus makes the semantics of communication simpler than ours as it does not rely on the distance between nodes; on the other hand it does not allow to enforce that a smart device should be either in a place or in another, but never in both. This can be easily represented in CaIT (again, see Proposition 3.1). CaIT has a finer control of inter-node communications as they depend on nodes' distance and transmission range of channels. Node mobility in CaIT is time constrained: in one time unit at most a fixed distance δ may be covered. Finally, Lanese et al.'s *end-user bisimilarity* shares most of the motivations of our bisimilarity. In the IoT-calculus, end users provide values to sensors and check actuators. They can also move nodes thus creating or removing connections, but they cannot observe channel communication. Thus, two systems with different connections between nodes are not end-user bisimilar. Unlike end-user bisimilarity, our notion of bisimilarity observes node mobility in a milder manner: the movements of a mobile node can be observed if the node either uses an actuator or transmits along a short-range channel or communicates its physical position. Finally, end-user bisimilarity is not preserved by parallel composition. Compositionality is recovered by observing also channel communication. The resulting bisimilarity models observation from the point of view of the other devices. Its distinguishing power is definitely stronger than that of our bisimilarity.

More recently, Bodei et al. [24,25] have proposed an untimed process calculus, IoT-LySA, supporting a control flow analysis that safely approximates the abstract behaviour of IoT systems. Essentially, they track how data spread from sensors to the logics of the network, and how physical data are manipulated. Intra-node generative communications in IoT-LySA are implemented through a shared store à la Linda [26]. In this manner, physical data are made available to software entities that analyse them and trigger the relevant actuators to perform the desired behaviour. The calculus adopts asynchronous multi-party communication among nodes taking care of node proximity (the topology is static). The dynamics of the calculus is given in terms of a reduction relation. No behavioural equivalences are defined.

Both the IoT-calculus, CaIT and IoT-LySA do not represent the physical processes which are part of a IoT system. Recently, we have proposed a calculus for cyber-physical processes, called CCPS [27,28], where the details of a physical process can be expressed in terms of difference equations. The calculus is equipped with a LTS semantics and a bisimulation-based behavioural semantics which supports compositional reasonings. The representation of the physical processes in CCPS is suitable for a formal study of cyber-physical attacks, i.e. attacks targeting physical devices (sensors and actuators) [29].

Our calculus takes inspiration from algebraic models for wireless systems [30–39,15,40]. All these models adopt broadcast communication on partial topologies, while we consider point-to-point communication, as in [8]. Our way of modelling network topology is taken from [30,32]. Paper [41] provides formal models for node mobility depending on the passage of time. Proposition 2.14 was inspired by [15]. A fully abstract observational theory for untimed ad hoc networks with broadcast communication can be found in [32,33]. Paper [15] contains a fully abstract observational theory for timed wireless systems with broadcast communication. Paper [38] provides a symbolic semantics for ad hoc networks.

Please cite this article in press as: R. Lanotte, M. Merro, A semantic theory of the Internet of Things, Inf. Comput. (2018), https://doi.org/10.1016/j.ic.2018.01.001

Vigo et al. [42] proposed a calculus for wireless-based cyber-physical (CPS) systems endowed with a theory that allows modelling and reasoning about cryptographic primitives, together with explicit notions of communication failure and unwanted communication. One of the main goal of the paper is a faithful representation of denial-of-service. However, as pointed out in [43], the calculus does not provide a notion of network topology, local broadcast and behavioural equivalence. It also lacks a clear distinction between physical components (sensors and actuators) and logical ones (processes). Compared to [42], paper [43] introduces a static network topology and enriches the theory with an harmony theorem.

As already said, CaIT has some similarities with the *synchronous languages* of the Esterel family [17,19,44]. In this setting, computations proceed in phases called "instants", which are quite similar to our time intervals. For instance, our timed reduction semantics has many points in common with Attar and Castellani's *CRL* synchronous reactive language [44], although *CRL* does not support mobility. The authors define two bisimulation equivalences. The first bisimulation formalises a *fine-grained* observation of programs: the observer is viewed as a program, which is able to interact with the observed program at any point of its execution. The second reflects a *coarse-grained* observation of programs: here the observer is viewed as part of the environment, which interacts with the observed program only at the start and the end of instants. The fine-grained bisimilarity is more in the style of a bisimulation for a process calculus.

CaIT is somehow reminiscent of De Nicola et al.'s SCEL language [45], a framework to model behaviour, knowledge, and data aggregation of Autonomic Systems.

Finally, the paper at hand extends the conference paper [1] in the following aspects: (i) full details of all proofs are spelled out, in particular we provide a non-standard proof of Theorem 5.4; (ii) the proof of Proposition 5.13 is a rare example of how compositional reasonings and up-to techniques can be used together to prove the bisimilarity of non-trivial systems; (iii) the design choices behind the primitives of the languages are explained in much more detail.

Acknowledgments

We thank Ilaria Castellani and Matthew Hennessy for their precious comments on an early draft. We thank Valentina Castiglioni for an early writing of the proof of the Harmony theorem. We thank the anonymous reviewers for their valuable comments and careful reviews.

Appendix A. Proofs

A.1. Proofs of Section 2

In order to prove Proposition 2.2 we need a technical lemma on time determinism of nodes.

Lemma A.1. *If* $n[P]_h^{\mu} \rightarrow_{\sigma} n'[P']_{h'}^{\mu'}$ *and* $n[P]_h^{\mu} \rightarrow_{\sigma} n''[P'']_{h''}^{\mu''}$ *then* $n = n' = n''$, $P' \equiv P''$, $\mu = \mu' = \mu''$ *and* $\mathrm{d}(h', h'') \leq 2\delta$.

Proof. By structural induction on P. □

Proof of Proposition 2.2. The proof is by rule induction on why $M \rightarrow_{\sigma} M'$. The most significant case is when $M \rightarrow_{\sigma} M'$ is derived by an application of rule (timemob):

$$\frac{n[\mathtt{I} \bowtie \prod_{i\in I} \lfloor \pi_i.P_i \rfloor Q_i | \prod_{j\in J} \sigma.P_j]_k^{\mathrm{m}} \not\rightarrow_{\tau} \quad \mathrm{d}(k,k') \leq \delta}{n[\mathtt{I} \bowtie \prod_{i\in I} \lfloor \pi_i.P_i \rfloor Q_i | \prod_{j\in J} \sigma.P_j]_k^{\mathrm{m}} \rightarrow_{\sigma} n[\mathtt{I} \bowtie \prod_{i\in I} Q_i | \prod_{j\in J} P_j]_{k'}^{\mathrm{m}}}$$

with $M = n[\mathtt{I} \bowtie \prod_{i\in I} \lfloor \pi_i.P_i \rfloor Q_i | \prod_{j\in J} \sigma.P_j]_k^{\mathrm{m}}$ and $M' = n[\mathtt{I} \bowtie \prod_{i\in I} Q_i | \prod_{j\in J} P_j]_{k'}^{\mathrm{m}}$. Suppose there exists M'' such that $M \rightarrow_{\sigma} M''$. Notice that, due to its structure, network M may perform a timed reduction only by an application of rule (timemob). Thus, by rule (timemob) we would have:

$$\frac{n[\mathtt{I} \bowtie \prod_{i\in I} \lfloor \pi_i.P_i \rfloor Q_i | \prod_{j\in J} \sigma.P_j]_k^{\mathrm{m}} \not\rightarrow_{\tau} \quad \mathrm{d}(k,k'') \leq \delta}{n[\mathtt{I} \bowtie \prod_{i\in I} \lfloor \pi_i.P_i \rfloor Q_i | \prod_{j\in J} \sigma.P_j]_k^{\mathrm{m}} \rightarrow_{\sigma} n[\mathtt{I} \bowtie \prod_{i\in I} Q_i'' | \prod_{j\in J} P_j'']_{k''}^{\mathrm{m}}}$$

with $M'' = n[\mathtt{I} \bowtie \prod_{i\in I} Q_i'' | \prod_{j\in J} P_j'']_{k''}^{\mathrm{m}}$. By Lemma A.1 it follows that $\prod_{i\in I} Q_i | \prod_{j\in J} P_j \equiv \prod_{i\in I} Q_i'' | \prod_{j\in J} P_j''$. Moreover, by triangular inequality it holds that $\mathrm{d}(k',k'') \leq \mathrm{d}(k,k') + \mathrm{d}(k,k'') \leq 2\delta$.

The other cases are similar. □

In order to prove maximal progress we need two simple lemmas whose proofs are omitted.

Lemma A.2. *If* $\prod_{i\in I} n_i[\mathtt{I}_i \bowtie P_i]_{h_i}^{\mu_i} \rightarrow_{\tau} M$ *then* $\prod_{i\in I} n_i[\mathtt{I}_i \bowtie P_i | Q_i]_{h_i}^{\mu_i} \rightarrow_{\tau} N$, *for some* N.

Lemma A.3. *If* $n[P]_h^{\mu} \not\rightarrow_{\sigma}$ *then for any process* Q *we have* $n[P|Q]_h^{\mu} \not\rightarrow_{\sigma}$.

Proof of Proposition 2.3. The proof is by rule induction on why $M \rightarrow_{\mathrm{i}} M'$. The most involved case is when $M \rightarrow_{\mathrm{i}} M'$ is derived by an application of rule (parp). This means that

$$M = \prod_{i\in I} n_i[\mathtt{I}_i \bowtie P_i | Q_i]^{\mu_i}_{h_i} \rightarrow_{\mathrm{i}} \prod_{i\in I} n_i[\mathtt{I}'_i \bowtie P'_i | Q_i]^{\mu_i}_{h'_i} = M'$$

because $\prod_{i\in I} n_i[\mathtt{I}_i \bowtie P_i]^{\mu_i}_{h_i} \rightarrow_{\mathrm{i}} \prod_{i\in I} n_i[\mathtt{I}'_i \bowtie P'_i]^{\mu_i}_{h'_i}$. By inductive hypothesis $\prod_{i\in I} n_i[\mathtt{I}_i \bowtie P_i]^{\mu_i}_{h_i} \not\rightarrow_\sigma$. We recall that rule (timepar) is the only one yielding timed reductions on parallel networks. Thus, if $\prod_{i\in I} n_i[\mathtt{I}_i \bowtie P_i]^{\mu_i}_{h_i} \not\rightarrow_\sigma$ it means that rule (timepar) could not be applied. Then, there are only two possibilities.

- Either $\prod_{i\in I} n_i[\mathtt{I}_i \bowtie P_i]^{\mu_i}_{h_i} \rightarrow_\tau N$, for some N. Then by Lemma A.2 we obtain $M \rightarrow_\tau N'$, for some N'. As rule (timepar) is the only rule yielding timed reductions from parallel networks, it follows that $M \not\rightarrow_\sigma$.
- Or $n_j[\mathtt{I}_j \bowtie P_j]^{\mu_j}_{h_j} \not\rightarrow_\sigma$, for some $j \in I$. Then by Lemma A.3 we have $n_j[\mathtt{I}_j \bowtie P_j | Q_j]^{\mu_j}_{h_j} \not\rightarrow_\sigma$. As rule (timepar) is the only rule yielding timed reductions from parallel networks, it follows that $M \not\rightarrow_\sigma$.

The remaining cases work smoothly. □

Proof of Proposition 2.4. The proof is by contradiction. We suppose there is no N such that $M \rightarrow_\sigma N$ and we prove that there is M' such that $M \rightarrow_{\mathrm{i}} M'$. We proceed by induction on the structure of M.

- Let $M = \mathbf{0}$. This case is not admissible because by an application or rule (timezero) we derive $M \rightarrow_\sigma M$.
- Let $M = n[P]^{\mu}_{h}$. As $M \not\rightarrow_\sigma$ and (timestat) and (timemob) are the only rules that could be used to derive a timed reduction from M, it follows that there are two possibilities.
  - Either $M \rightarrow_\tau M'$, for some M', and we are done.
  - Or P has not the proper structure for applying rule (timestat) or rule (timemob). This means that $P \equiv P_1 | P_2$, with $P_1 = \rho.P'_1$ and $\rho \in \{@(x), s?(x), a!v\}$. In this case, by an application of one among the rules (pos), (sensread), (actunchg), and (actchg), followed by an application of rule (parp), we can infer $M \rightarrow_{\mathrm{i}} M'$, for some M'.
- Let $M = M_1 | M_2$, for some M_1 and M_2. As $M \not\rightarrow_\sigma$ and (timepar) is the only rule which could be used to derive a timed reduction from M, it follows that there are two possibilities.
  - Either $M \rightarrow_\tau M'$, for some M'; hence $M \rightarrow_{\mathrm{i}} M'$ and we are done.
  - Or at least one among M_1 and M_2 cannot perform a timed reduction. Suppose $M_1 \not\rightarrow_\sigma$; by inductive hypothesis there is M'_1 such that $M_1 \rightarrow_{\mathrm{i}} M'_1$. By an application of rule (parn) we derive $M \rightarrow_{\mathrm{i}} M'_1 | M_2$.
- Let $M = (\boldsymbol{\nu} c)M_1$. This case requires an easy application of the inductive hypothesis. □

In order to prove Proposition 2.5, we need a couple of technical definitions and lemmas.

Definition A.4. Let use define pfxi() as the function that given a process P returns an upper bound to the number of the *untimed prefixes* that can give rise to an instantaneous reduction when P is plugged in a node.

$$\begin{array}{lll} \mathrm{pfxi}(\mathsf{nil}) = 0 & \mathrm{pfxi}(\sigma.P) = 0 & \mathrm{pfxi}(\rho.P) = 1 + \mathrm{pfxi}(P) \ (\text{if } \rho \neq \sigma) \\ \mathrm{pfxi}(X) = \infty & \mathrm{pfxi}(\mathsf{fix}\, X.P) = \mathrm{pfxi}(P) & \mathrm{pfxi}(\lfloor \pi.P \rfloor Q) = 1 + \mathrm{pfxi}(P) \\ \mathrm{pfxi}([b]P;\, Q) = \max(\mathrm{pfxi}(P), \mathrm{pfxi}(Q)) & & \mathrm{pfxi}(P \mid Q) = \mathrm{pfxi}(P) + \mathrm{pfxi}(Q). \end{array}$$

Lemma A.5. *For any closed process P,* $\mathrm{pfxi}(P)$ *is finite.*

Proof. The proof is by structural induction on P. The only interesting case is when $P = \mathsf{fix}\, X.P_1$, as P_1 may contain the process variable X and $\mathrm{pfxi}(X) = \infty$. However, in our calculus we only admit time-guarded recursion. Thus, X may occur in P_1 only if guarded by at least one σ prefix, and $\mathrm{pfxi}(\sigma.Q) = 0$, for any Q. It follows that $\mathrm{pfxi}(\mathsf{fix}\, X.P_1) \in \mathbb{N}$, for any P_1. □

Thus, the function redi() of Definition 5.2 provides an upper bound to the number of consecutive instantaneous reductions performed by a system. From Lemma A.5 it can be easily proved that $\mathrm{redi}(M)$ is always finite.

Lemma A.6. *For any network M,* $\mathrm{redi}(M)$ *is finite.*

Proof. By induction on the structure of M. □

Proof of Proposition 2.5. By induction on the structure of M; it follows from Lemma A.6. □

Please cite this article in press as: R. Lanotte, M. Merro, A semantic theory of the Internet of Things, Inf. Comput. (2018), https://doi.org/10.1016/j.ic.2018.01.001

A.2. Proofs of Section 3

Let us prove Proposition 3.1. For that we need a technical lemma.

Lemma A.7. *If* $Sys\ (\rightarrow_{\mathrm{i}}^{*}\ \rightarrow_{\sigma})^{*}\ Sys'$ *then* $Sys' \equiv Phone'|Home'$ *where:*

- $Phone' = n_P[\mathtt{I}_P \bowtie BoilerCtrl \mid LightCtrl]^{\mathrm{m}}_{l'}$, *for some* l', *with* $\mathtt{I}_P(mode) = \mathsf{auto}$
- $Home' = LR_1|LR_2|BoilerMng$, *for some* LR_1 *and* LR_2,
- $BoilerMng = n_B[\mathtt{I}_B \bowtie BoilerMng]^{\mathrm{s}}_{loc2}$, *with* $\mathtt{I}_B(temp) = \Theta$.

Proof. The proof is by mathematical induction on the integer j such that $Sys\ (\rightarrow_{\mathrm{i}}^{*}\ \rightarrow_{\sigma})^{j}\ Sys'$.

The case $j = 0$ is trivial.

Let us move on the inductive case. Let $Sys\ (\rightarrow_{\mathrm{i}}^{*}\ \rightarrow_{\sigma})^{j}\ Sys_1$, for $j > 0$. By inductive hypothesis we have: $Sys_1 \equiv Phone_1|Home_1$ where:

- $Phone_1 = n_P[\mathtt{I}_P \bowtie BoilerCtrl \mid LightCtrl]^{\mathrm{m}}_{l_1}$, for some l_1, with $\mathtt{I}_P(mode) = \mathsf{auto}$
- $Home_1 = LR_1|LR_2|BoilerMng$, for some LR_1 and LR_2,
- $BoilerMng = n_B[\mathtt{I}_B \bowtie BoilerMng]^{\mathrm{s}}_{loc2}$, with $\mathtt{I}_B(temp) = \Theta$.

We recall that sensor changes are not modelled in the reduction semantics as they require the intervention of the physical environment. Thus, the value of all sensors will remain unchanged during the reduction sequence. We show that whenever $Sys_1 \rightarrow_{\mathrm{i}}^{*} \rightarrow_{\sigma} Sys'$, then Sys' has still the same structure as Sys_1. Let us consider a portion of Sys_1 composed by the phone and the boiler manager. Then, we have the following sequence of instantaneous reductions. We recall that $\mathtt{I}_P(mode) = \mathsf{auto}$ and $\mathtt{I}_B(temp) = \Theta$.

$$\begin{aligned}
Phone_1 \mid BoilerMng \;&=\; n_P[\mathtt{I}_P \bowtie BoilerCtrl \mid LightCtrl]^{\mathrm{m}}_{l_1} \;\mid\; n_B[\mathtt{I}_B \bowtie BoilerMng]^{\mathrm{s}}_{loc2}\\
&\rightarrow_{\mathrm{i}} n_P[\mathtt{I}_P \bowtie \lfloor \overline{b}\langle \mathsf{auto}\rangle.\sigma.BoilerCtrl \rfloor \ldots \mid LightCtrl]^{\mathrm{m}}_{l_1} \;\mid\; n_B[\mathtt{I}_B \bowtie BoilerMng]^{\mathrm{s}}_{loc2}\\
&\rightarrow_{\mathrm{i}} n_P[\mathtt{I}_P \bowtie \sigma.BoilerCtrl \mid LightCtrl]^{\mathrm{m}}_{l_1} \;\mid\; n_B[\mathtt{I}_B \bowtie TempCtrl]^{\mathrm{s}}_{loc2}\\
&\rightarrow_{\mathrm{i}} n_P[\mathtt{I}_P \bowtie \sigma.BoilerCtrl \mid LightCtrl]^{\mathrm{m}}_{l_1} \;\mid\; n_B[\mathtt{I}_B \bowtie boiler!\mathsf{off}.\sigma.BoilerMng]^{\mathrm{s}}_{loc2}\\
&\rightarrow_{\mathrm{i}} n_P[\mathtt{I}_P \bowtie \sigma.BoilerCtrl \mid LightCtrl]^{\mathrm{m}}_{l_1} \;\mid\; n_B[\mathtt{I}_B[boiler \mapsto \mathsf{off}] \bowtie \sigma.BoilerMng]^{\mathrm{s}}_{loc2}\\
&=\; n_P[\mathtt{I}_P \bowtie \sigma.BoilerCtrl \mid LightCtrl]^{\mathrm{m}}_{l_1} \;\mid\; n_B[\mathtt{I}_B \bowtie \sigma.BoilerMng]^{\mathrm{s}}_{loc2}.
\end{aligned}$$

Now, both the phone and the boiler manager can only perform a timed reduction. However, the whole system may have further instantaneous reductions depending whether the phone is in position to interact with the light managers of the house. In any case, thanks to (i) *well-timedness* (Proposition 2.5), (ii) *patience* (Proposition 2.4), (iii) rule (parn), (iv) rule (struct) we will eventually have a reduction sequence of the form:

$$Phone_1 \mid LR_1 \mid LR_2 \mid BoilerMng \;\rightarrow_{\mathrm{i}}^{*}\; \rightarrow_{\sigma}\; Phone' \mid LR'_1 \mid LR'_2 \mid BoilerMng$$

in which $\mathtt{I}_P(mode) = \mathsf{auto}$ and $\mathtt{I}_B(temp) = \Theta$ (the reduction semantics cannot change sensor values) and $Phone'$ is exactly as $Phone_1$ except for the fact that is located at a possibly new location l', with $\mathrm{d}(l_1, l') \leq 1$. □

Proof of Proposition 3.1. By Lemma A.7 we deduce that Sys' preserves the structure of Sys and also the value of its sensors. Let us prove the four cases of the proposition, one by one.

1. Let us consider the evolution of $Sys'[mode \mapsto \mathsf{man}]$. By inspection on the definitions in Table 4 it is easy to derive that

$$\begin{aligned}
Sys'[mode \mapsto \mathsf{man}] \;&\equiv\; Phone'[mode \mapsto \mathsf{man}]|LR_1|LR_2|n_B[\mathtt{I}_B \bowtie BoilerMng]^{\mathrm{s}}_{loc2}\\
&\rightarrow_{\mathrm{i}}^{*} \rightarrow_{\sigma}\; Phone'[mode \mapsto \mathsf{man}]|LR'_1|LR'_2|n_B[\mathtt{I}'_B \bowtie BoilerManual]^{\mathrm{s}}_{loc2}
\end{aligned}$$

   with $\mathtt{I}'_B(boiler) = \mathsf{on}$.
2. Let us consider the evolution of $Sys'[temp \mapsto t]$, with $t < \Theta$. We spell out this case in more detail. We recall that the sensor *mode* of the phone is set to auto. We also recall that, by an application of rules (parn) and (struct), if a parallel component can execute an instantaneous reduction then the whole network can execute the same reduction. Thus, in the following we concentrate on the reductions deriving from the phone and the boiler manager when $t < \Theta$.

Please cite this article in press as: R. Lanotte, M. Merro, A semantic theory of the Internet of Things, Inf. Comput. (2018), https://doi.org/10.1016/j.ic.2018.01.001

$$
\begin{aligned}
(\mathit{Phone}' \mid \mathit{BoilerMng})[\mathit{temp}\mapsto t] & \\
&= n_P[\mathtt{I}_P \bowtie \mathit{BoilerCtrl} \mid \mathit{LightCtrl}]^{\mathrm{m}}_{l'} \mid n_B[\mathtt{I}_B[\mathit{temp} \mapsto t] \bowtie \mathit{BoilerMng}]^{\mathrm{s}}_{loc2} \\
&\to_{\mathrm{i}} n_P[\mathtt{I}_P \bowtie \lfloor \overline{b}\langle \mathsf{auto}\rangle.\sigma.\mathit{BoilerCtrl}\rfloor \ldots \mid \mathit{LightCtrl}]^{\mathrm{m}}_{l'} \mid n_B[\mathtt{I}'_B \bowtie \mathit{BoilerMng}]^{\mathrm{s}}_{loc2} \\
&\to_{\mathrm{i}} n_P[\mathtt{I}_P \bowtie \sigma.\mathit{BoilerCtrl} \mid \mathit{LightCtrl}]^{\mathrm{m}}_{l'} \mid n_B[\mathtt{I}'_B \bowtie \mathit{TempCtrl}]^{\mathrm{s}}_{loc2} \\
&\to_{\mathrm{i}} n_P[\mathtt{I}_P \bowtie \sigma.\mathit{BoilerCtrl} \mid \mathit{LightCtrl}]^{\mathrm{m}}_{l'} \mid n_B[\mathtt{I}'_B \bowtie \mathit{boiler}!\mathsf{on}.\sigma.\mathit{BoilerMng}]^{\mathrm{s}}_{loc2} \\
&\to_{\mathrm{i}} n_P[\mathtt{I}_P \bowtie \sigma.\mathit{BoilerCtrl} \mid \mathit{LightCtrl}]^{\mathrm{m}}_{l'} \mid n_B[\mathtt{I}'_B[\mathit{boiler} \mapsto \mathsf{on}] \bowtie \sigma.\mathit{BoilerMng}]^{\mathrm{s}}_{loc2} \\
&= n_P[\mathtt{I}_P \bowtie \sigma.\mathit{BoilerCtrl} \mid \mathit{LightCtrl}]^{\mathrm{m}}_{l'} \mid n_B[\mathtt{I}''_B \bowtie \sigma.\mathit{BoilerMng}]^{\mathrm{s}}_{loc2}
\end{aligned}
$$

where $\mathtt{I}''_B(\mathit{temp}) = t < \Theta$ and $\mathtt{I}''_B(\mathit{boiler}) = \mathsf{on}$. Now, both the phone and the boiler manager can only perform a timed reduction. However, the whole system may have further instantaneous reductions depending whether the phone is in position to interact with the light managers of the house. In any case, thanks to *well-timedness* (Proposition 2.5) and *patience* (Proposition 2.4) we will eventually have a reduction sequence of the form:

$$\mathit{Sys}'[\mathit{temp} \mapsto t] \to_{\mathrm{i}}^* \to_{\sigma} \mathit{Phone}''|LR'_1|LR'_2|n_B[\mathtt{I}''_B \bowtie \mathit{BoilerMng}]^{\mathrm{s}}_{loc2}$$

where $\mathtt{I}''_B(\mathit{temp}) = t$, $\mathtt{I}''_B(\mathit{boiler}) = \mathsf{on}$, and the mobile phone may have moved to a new location l'', with $\mathrm{d}(l', l'') = 1$.

3. Let us consider the evolution of $\mathit{Sys}'[\mathit{temp} \mapsto t]$, with $t \geq \Theta$. Here, similarly to the previous case, we can derive:

$$
\begin{aligned}
\mathit{Sys}'[\mathit{temp} \mapsto t] \quad &\equiv \quad \mathit{Phone}'|LR_1|LR_2|n_B[\mathtt{I}_B[\mathit{temp} \mapsto t] \bowtie \mathit{BoilerMng}]^{\mathrm{s}}_{loc2} \\
&\to_{\mathrm{i}}^* \to_{\sigma} \mathit{Phone}''|LR'_1|LR'_2|n_B[\mathtt{I}'_B \bowtie \mathit{BoilerMng}]^{\mathrm{s}}_{loc2}
\end{aligned}
$$

with $\mathtt{I}'_B(\mathit{temp}) = t$ and $\mathtt{I}'_B(\mathit{boiler}) = \mathsf{off}$.

4. We prove only the implication from left to right. The other is similar. We know that $\mathit{Sys}' \to_{\mathrm{i}}^* \mathit{Sys}'' \downarrow_{\mathit{light}_1!\mathsf{on}}$. By Lemma A.7 we know the structure of Sys'. We recall that in Sys the actuator light_1 is set to off. Notice also that this actuator is exclusively managed via the process $\mathit{LightMng}_1$, running at the stationary node n_1, located al $\mathit{loc}1$. More precisely, the actuator light_1 can be modified by $\mathit{LightMng}_1$ only after a synchronisation at the short-range channel c_1. We recall that $\mathrm{rng}(c_1) = 0$. We also recall that mobile nodes can change their location only by executing a timed reduction via rule (timemob). We fixed $\delta = 1$, which is the maximum distance that a mobile node can afford within a time unit. Thus, if $\mathit{Sys}' \to_{\mathrm{i}}^* \mathit{Sys}'' \downarrow_{\mathit{light}_1!\mathsf{on}}$ there are two possibilities:
   - either the mobile phone in current time interval is located at $\mathit{loc}1$;
   - or the mobile phone was located at $\mathit{loc}1$ in the previous time interval, and in the current time interval it is at a location l', with $\mathrm{d}(l', \mathit{loc}1) = 1$, as $\delta = 1$.

   In the first case, the light manager $\mathit{LightMng}_2$, located at $\mathit{loc}4$, has necessarily set the actuator light_2 to off. This is because $\mathrm{rng}(c_2) = 0$, $\mathrm{d}(\mathit{loc}1, \mathit{loc}4) = 3$, and the only manner to switch on light_2 is to place the mobile phone at $\mathit{loc}4$. However, as the mobile phone is currently at $\mathit{loc}1$, and $\delta = 1$, this could have happened only 3 time instants ago. By that time, the timeout in $\mathit{LightMng}_2$ has already switched off the light.
   The second case, when the mobile phone is currently located at some location l', with $\mathrm{d}(l', \mathit{loc}1) = 1$, is similar. This is because $\mathrm{d}(\mathit{loc}1, \mathit{loc}4) = 3$, and by triangular inequality $\mathrm{d}(l', \mathit{loc}4) \geq 2$. Thus, the phone is far enough to ensure that timeout of $\mathit{LightMng}_2$ already fired to switch off light_2. □

A.3. Proofs of Section 4

This section is devoted to the proof of the *Harmony Theorem* (Theorem 4.1). We start with a technical lemma that provides the structure of a process depending on its possible actions.

Lemma A.8. *Let P be a process.*

1. *If $P \xrightarrow{\sigma} P'$ then $P \equiv \prod_{i\in I} \lfloor \pi_i.P_i \rfloor Q_i \mid \prod_{j\in J} \sigma.P_j$ and $P' \equiv \prod_{i\in I} Q_i \mid \prod_{j\in J} P_j$, for appropriate index sets, prefixes and processes.*
2. *If $P \xrightarrow{s?v} P'$ then there are P_1 and Q such that $P \equiv s?(x).P_1 \mid Q$ and $P' \equiv P_1\{^v/_x\} \mid Q$.*
3. *If $P \xrightarrow{a!v} P'$ then there are P_1 and Q such that $P \equiv a!v.P_1 \mid Q$ and $P' \equiv P_1 \mid Q$.*
4. *If $P \xrightarrow{@h} P'$ then there are P_1 and Q such that $P \equiv @(x).P_1 \mid Q$ and $P' \equiv P_1\{^h/_x\} \mid Q$.*
5. *If $P \xrightarrow{\overline{c}v} P'$ then there are P_1, Q_1 and Q such that $P \equiv \lfloor \overline{c}\langle v\rangle.P_1 \rfloor Q_1 \mid Q$ and $P' \equiv P_1 \mid Q$.*
6. *If $P \xrightarrow{cv} P'$, then there exist P_1, Q_1, Q s.t. $P \equiv \lfloor c(x).P_1 \rfloor Q_1 \mid Q$ and $P' \equiv P_1\{^v/_x\} \mid Q$.*
7. *If $P \xrightarrow{\tau} P'$ then there are P_1, P_2, Q_1, Q_2, R, and c with $\mathrm{rng}(c) = -1$, such that $P \equiv \lfloor c(x).P_1 \rfloor Q_1 \mid \lfloor \overline{c}\langle v\rangle.P_2 \rfloor Q_2 \mid R$ and $P' \equiv P_1\{^v/_x\} \mid P_2 \mid R$.*

Please cite this article in press as: R. Lanotte, M. Merro, A semantic theory of the Internet of Things, Inf. Comput. (2018), https://doi.org/10.1016/j.ic.2018.01.001

Proof. Let us start with item (1). We proceed by rule induction on why $P \xrightarrow{\sigma} P'$.

- Let $P \xrightarrow{\sigma} P'$ by an application of rule (TimeNil); then the result is immediate for $I = J = \emptyset$.
- Let $P \xrightarrow{\sigma} P'$ by an application of rule (Delay) with $P = \sigma.P_1$ and $P' = P_1$. Thus, for $I = \emptyset$ and $J = \{1\}$ we have $P = \sigma.P_1 \equiv \mathsf{nil} \mid \sigma.P_1$ and $P' = P_1 \equiv \mathsf{nil}|P_1$.
- Let $P \xrightarrow{\sigma} P'$ by an application of rule (Timeout) with $P = \lfloor \pi_1.P_1 \rfloor Q_1$ and $P' = Q_1$. Thus, for $I = \{1\}$ and $J = \emptyset$ we have $P = \lfloor \pi_1.P_1 \rfloor Q_1 \equiv \lfloor \pi_1.P_1 \rfloor Q_1|\mathsf{nil}$ and $P' = Q_1 \equiv Q_1|\mathsf{nil}$.
- Let $P \xrightarrow{\sigma} P'$ by an application of rule (TimeParP) with $P = R_1 \mid R_2$ and $P' = R_1' \mid R_2'$, because $R_1 \xrightarrow{\sigma} R_1'$ and $R_2 \xrightarrow{\sigma} R_2'$. By inductive hypothesis, there exist I, J, I' and J' such that $R_1 \equiv \prod_{i\in I} \lfloor \pi_i.P_i \rfloor Q_i \mid \prod_{j\in J} \sigma.P_j$, $R_1' \equiv \prod_{i\in I} Q_i \mid \prod_{j\in J} P_j$, $R_2 \equiv \prod_{i'\in I'} \lfloor \pi_{i'}.P_{i'} \rfloor Q_{i'} \mid \prod_{j'\in J'} \sigma.P_{j'}$ and $R_2' \equiv \prod_{i'\in I'} Q_{i'} \mid \prod_{j'\in J'} P_{j'}$. To conclude this case we choose as index sets $\bar{I} = I \cup I'$ and $\bar{J} = J \cup J'$.
- Let $P \xrightarrow{\sigma} P'$ by an application of rule (Fix) with $P = \mathsf{fix}\, X.P_1$ and $P' = P_2$, because $P_1\{^{\mathsf{fix}\, X.P_1}/_X\} \xrightarrow{\sigma} P_2$. By inductive hypothesis, there exist I and J such that $P_1\{^{\mathsf{fix}\, X.P_1}/_X\} \equiv \prod_{i\in I} \lfloor \pi_i.P_i \rfloor Q_i \mid \prod_{j\in J} \sigma.P_j$ and $P_2 \equiv \prod_{i\in I} Q_i \mid \prod_{j\in J} P_j$. By structural congruence we have $\mathsf{fix}\, X.P_1 \equiv P_1\{^{\mathsf{fix}\, X.P_1}/_X\}$ and therefore $P = \mathsf{fix}\, X.P_1 \equiv \prod_{i\in I} \lfloor \pi_i.P_i \rfloor Q_i \mid \prod_{j\in J} \sigma.P_j$.

Let us prove the item (2) of the lemma. We proceed by rule induction on why $P \xrightarrow{s?v} P'$.

- Let $P \xrightarrow{s?v} P'$ by an application of rule (Sensor) with $P = s?(x).P_1$ and $P' = P_1\{^v/_x\}$. This case is easy.
- Let $P \xrightarrow{s?v} P'$ by an application of rule (Fix) with $P = \mathsf{fix}\, X.P_1$ and $P' = P_2$, because $P_1\{^{\mathsf{fix}\, X.P_1}/_X\} \xrightarrow{s?v} P_2$. By inductive hypothesis there exist P_3 and Q_1 such that $P_1\{^{\mathsf{fix}\, X.P_1}/_X\} \equiv s?(x).P_3 \mid Q_1$ and $P_2 \equiv P_3\{^v/_x\} \mid Q_1$. By structural congruence $P = \mathsf{fix}\, X.P_1 \equiv P_1\{^{\mathsf{fix}\, X.P_1}/_X\} \equiv s?(x).P_3 \mid Q_1$.
- Let $P \xrightarrow{s?v} P'$ by an application of rule (ParP) with $P = P_1 \mid R$ and $P' = P_1' \mid R$, because $P_1 \xrightarrow{s?v} P_1'$. By inductive hypothesis there exist P_2 and Q_1 such that $P_1 \equiv s?(x).P_2 \mid Q_1$ and $P_1' \equiv P_2\{^v/_x\} \mid Q_1$, thus, the thesis holds for $Q = Q_1 \mid R$.

The cases (3), (4), (5) and (6) are analogous to previous items. We prove (7). We proceed by rule induction on why $P \xrightarrow{\tau} P'$.

- Let $P \xrightarrow{\tau} P'$ by an application of rule (Com). Then our result follows by application of the items (5) and (6) of the proposition. We need to work up to structural congruence.
- Let $P \xrightarrow{\tau} P'$ by an application of rule (ParP) or (Fix). This case is analogous to the corresponding ones in (2). □

The following lemma is similar to the previous one but it deals with networks rather than processes.

Lemma A.9.

1. *If* $M \xrightarrow{\overline{c}v@h} M'$ *then* $M \equiv (\nu\tilde{g})\big(n[\mathtt{I} \bowtie \lfloor \overline{c}\langle v\rangle.P \rfloor P' \mid Q]_h^\mu | N\big)$ *and* $M' \equiv (\nu\tilde{g})\big(n[\mathtt{I} \bowtie P \mid Q]_h^\mu | N\big)$, *for some* $n, P, P', Q, \mu, N, \tilde{g}$, *with* $c \notin \tilde{g}$.
2. *If* $M \xrightarrow{cv@h} M'$ *then* $M \equiv (\nu\tilde{g})\big(n[\mathtt{I} \bowtie \lfloor c(x).P \rfloor P' \mid Q]_h^\mu | N\big)$ *and* $M' \equiv (\nu\tilde{g})\big(n[\mathtt{I} \bowtie P\{^v/_x\} \mid Q]_h^\mu | N\big)$, *for some* $n, P, P', Q, \mu, N, \tilde{g}$, *with* $c \notin \tilde{g}$.

Proof. We start with the proof of item (1). We proceed by rule induction.

- Let $M \xrightarrow{\overline{c}v@k} M'$ by an application of rule (SndN) because $P \xrightarrow{\overline{c}v} P'$ and $\mathrm{rng}(c) \geq 0$, with $M = n[\mathtt{I} \bowtie P]_k^\mu$ and $M' = n[\mathtt{I} \bowtie P']_k^\mu$. Lemma A.8(5) ensures that since $P \xrightarrow{\overline{c}v} P'$ then there exist P_1, Q_1, Q such that $P \equiv \lfloor \overline{c}\langle v\rangle.P_1 \rfloor Q_1 \mid Q$ and $P' \equiv P_1 \mid Q$. This implies $M \equiv n[\mathtt{I} \bowtie \lfloor \overline{c}\langle v\rangle.P_1 \rfloor Q_1 \mid Q]_k^\mu$ and $M' \equiv n[\mathtt{I} \bowtie P_1 \mid Q]_k^\mu$.
- Let $M \xrightarrow{\overline{c}v@k} M'$ by an application of rule (ParN) because $M_1 \xrightarrow{\overline{c}v@k} M_1'$, with $M = M_1|M_2$ and $M' = M_1'|M_2$. By inductive hypothesis, since $M_1 \xrightarrow{\overline{c}v@k} M_1'$, there exist $n, P_1, P_1', Q_1, \mu, k, N_1, \tilde{g}$ such that $c \notin \tilde{g}$ and

$$M_1 \equiv (\nu\tilde{g})n[\mathtt{I} \bowtie \lfloor \overline{c}\langle v\rangle P_1 \rfloor P_1' \mid Q_1]_k^\mu | N_1 \quad \text{and} \quad M_1' \equiv (\nu\tilde{g})n[\mathtt{I} \bowtie P_1 \mid Q_1]_k^\mu | N_1.$$

Hence $M \equiv (\nu\tilde{g})n[\mathtt{I} \bowtie \lfloor \overline{c}\langle v\rangle P_1 \rfloor P_1' \mid Q_1]_k^\mu | N_1 | M_2$ and the system M' is such that $M' \equiv (\nu\tilde{g})n[\mathtt{I} \bowtie P_1 \mid Q_1]_k^\mu | N_1 | M_2$. This concludes the case, for $N = N_1|M_2$.

Please cite this article in press as: R. Lanotte, M. Merro, A semantic theory of the Internet of Things, Inf. Comput. (2018), https://doi.org/10.1016/j.ic.2018.01.001

- Let $M \xrightarrow{\overline{c}v@k} M'$ by an application of rule (Res) because $M_1 \xrightarrow{\overline{c}v@k} M'_1$, with $c \neq c'$, $M = (\boldsymbol{\nu} c')M_1$ and $M' = (\boldsymbol{\nu} c')M'_1$. By inductive hypothesis, there exist $n, P_1, P'_1, Q_1, \mu, k, N_1, \tilde{g}$ such that $c \notin \tilde{g}$ and $M_1 \equiv (\boldsymbol{\nu}\tilde{g})n[\mathtt{I} \bowtie \lfloor\overline{c}\langle v\rangle P_1 \rfloor P'_1 \mid Q_1]^{\mu}_k | N_1$ and $M'_1 \equiv (\boldsymbol{\nu}\tilde{g})n[\mathtt{I} \bowtie P_1 \mid Q_1]^{\mu}_k | N_1$. Hence $M \equiv (\boldsymbol{\nu} c')(\boldsymbol{\nu}\tilde{g})n[\mathtt{I} \bowtie \lfloor\overline{c}\langle v\rangle P_1 \rfloor P'_1 \mid Q_1]^{\mu}_k | N_1$ and $M' \equiv (\boldsymbol{\nu} c')(\boldsymbol{\nu}\tilde{g})n[\mathtt{I} \bowtie P_1 \mid Q_1]^{\mu}_k | N_1$. Thus, since $c \notin (\boldsymbol{\nu} c')(\boldsymbol{\nu}\tilde{g})$ this concludes the case.

The remaining item (2) is analogous by applying Lemma A.8(6). □

The following lemma says that structural congruence is a (strong) bisimulation.

Lemma A.10. *If $M \xrightarrow{\alpha} M'$ and $M \equiv N$ then there is N' such that $N \xrightarrow{\alpha} N'$ and $M' \equiv N'$.*

Proof of Theorem 4.1. We have to prove the following sub-results:

1. If $M \xrightarrow{\tau} M'$ then $M \rightarrow_{\tau} M'$.
2. If $M \rightarrow_{\tau} M'$ then $M \xrightarrow{\tau} \equiv M'$.
3. If $M \xrightarrow{a} M'$ then $M \rightarrow_{a} M'$.
4. If $M \rightarrow_{a} M'$ then $M \xrightarrow{a} \equiv M'$.
5. If $M \xrightarrow{\sigma} M'$ then $M \rightarrow_{\sigma} M'$.
6. If $M \rightarrow_{\sigma} M'$ then $M \xrightarrow{\sigma} \equiv M'$.

Let us start with the sub-result (1). The proof is by rule induction on why $M \xrightarrow{\tau} M'$.

- Let $M \xrightarrow{\tau} M'$ by an application of rule (SensRead), with $M = n[\mathtt{I} \bowtie P]^{\mu}_h$ and $M' = n[\mathtt{I} \bowtie P']^{\mu}_h$, because $\mathtt{I}(s) = v$ and $P \xrightarrow{s?v} P'$. By Lemma A.8(2) there exist P_1, Q such that $P \equiv s?(x).P_1 \mid Q$ and $P' \equiv P_1\{^v/_x\} \mid Q$. Then we can apply the reduction rules (sensread) and (parp) to infer $M \rightarrow_{\tau} M'$, as required.
- Let $M \xrightarrow{\tau} M'$ by an application of rule (Pos). This case follows by an application of Lemma A.8(4) together with reduction rules (pos) and (parp).
- Let $M \xrightarrow{\tau} M'$ by an application of rule (LocCom), with $M = n[\mathtt{I} \bowtie P]^{\mu}_k$ and $M' = n[\mathtt{I} \bowtie P']^{\mu}_k$, because $P \xrightarrow{\tau} P'$. By Lemma A.8(7), $P \xrightarrow{\tau} P'$ ensures that there exist P_1, P_2, Q_1, Q_2, R, c with $\mathrm{rng}(c) = -1$ such that $P \equiv \lfloor c(x).P_1 \rfloor Q_1 \mid \lfloor\overline{c}\langle v\rangle.P_2 \rfloor Q_2 \mid R$ and $P' \equiv P_1\{^v/_x\} \mid P_2 \mid R$. By structural congruence we have

  $$M \equiv n[\mathtt{I} \bowtie \lfloor c(x).P_1 \rfloor Q_1 \mid \lfloor\overline{c}\langle v\rangle.P_2 \rfloor Q_2 \mid R]^{\mu}_k$$

  and analogously $M' \equiv n[\mathtt{I} \bowtie P_1\{^v/_x\} \mid P_2 \mid R]^{\mu}_k$. Hence, by an application of rules (struct) and (loccom) we can infer $M \rightarrow_{\tau} M'$.
- Let $M \xrightarrow{\tau} M'$ by an application of rule (ActUnChg). This case follows by an application of Lemma A.8(3) together with an application of reduction rules (actunchg) and (parp).
- Let $M \xrightarrow{\tau} M'$ by an application of rule (ParN), with $M = M_1 | M_2$ and $M' = M'_1 | M_2$, because $M_1 \xrightarrow{\tau} M'_1$. By inductive hypothesis $M_1 \rightarrow_{\tau} M'_1$. Therefore, by an application of rule (parn) we can infer $M \rightarrow_{\tau} M'$.
- Let $M \xrightarrow{\tau} M'$ by an application of rule (Res), with $M = (\boldsymbol{\nu}\tilde{g})M_1$ and $M' = (\boldsymbol{\nu}\tilde{g})M'_1$, because $M_1 \xrightarrow{\tau} M'_1$. By inductive hypothesis $M_1 \rightarrow_{\tau} M'_1$. Therefore, by an application of the reduction rule (res) we derive $M \rightarrow_{\tau} M'$.
- Let $M \xrightarrow{\tau} M'$ by an application of rule (GlbCom), with $M = M_1 | M_2$ and $M' = M'_1 | M'_2$, because $M_1 \xrightarrow{\overline{c}v@h} M'_1$ and $M_2 \xrightarrow{cv@k} M'_2$ and $\mathrm{d}(h,k) \leq \mathrm{rng}(c)$. Since $M_1 \xrightarrow{\overline{c}v@h} M'_1$, Lemma A.9(1) guarantees that $M_1 \equiv (\boldsymbol{\nu}\tilde{g})n[\mathtt{I} \bowtie \lfloor\overline{c}\langle v\rangle P \rfloor P' \mid R]^{\mu}_h | N$ and $M'_1 \equiv (\boldsymbol{\nu}\tilde{g})n[\mathtt{I} \bowtie P \mid R]^{\mu}_h | N$, for some $n, P, P', R, \mu, h, N, \tilde{g}$. Furthermore, by Lemma A.9(2) there exist $m, Q, Q', R', \mu, k, N', \tilde{g}'$ such that $M_2 \equiv (\boldsymbol{\nu}\tilde{g}')m[\mathtt{I} \bowtie \lfloor c(x)Q \rfloor Q' \mid R']^{\mu}_k | N'$ and $M'_2 \equiv (\boldsymbol{\nu}\tilde{g}')m[\mathtt{I} \bowtie Q\{^v/_x\} \mid R']^{\mu}_k | N'$. Therefore, by applying the reduction rules (struct), (res), (glbcom), (parp) and (parn) we can infer $M \rightarrow_{\tau} M'$.

Let us prove the sub-result (2) by rule induction on why $M \rightarrow_{\tau} M'$.

- Let $M \rightarrow_{\tau} M'$ by rule (sensread), with $M = n[\mathtt{I} \bowtie s?(x).P \mid Q]^{\mu}_h$ and $M' = n[\mathtt{I} \bowtie P\{^v/_x\} \mid Q]^{\mu}_h$, because $\mathtt{I}(s) = v$. Hence, by rule (Sensor) we have $s?(x).P \xrightarrow{s?v} P\{^v/_x\}$, by rule (ParP) we have $s?(x).P \mid Q \xrightarrow{s?v} P\{^v/_x\} \mid Q$ and finally by rule (SensRead) we have $n[\mathtt{I} \bowtie s?(x).P \mid Q]^{\mu}_h \xrightarrow{\tau} n[\mathtt{I} \bowtie P\{^v/_x\} \mid Q]^{\mu}_h$.
- Let $M \rightarrow_{\tau} M'$ by applying rule (pos), with $M = n[\mathtt{I} \bowtie @(x).P \mid Q]^{\mu}_h$ and $M' = n[\mathtt{I} \bowtie P\{^x/_h\} \mid Q]^{\mu}_h$. We get $M \xrightarrow{\tau} \equiv M'$ by applying rules (PosP), (ParP) and (Pos).
- Let $M \rightarrow_{\tau} M'$ by an application of rule (actunchg). This case is similar to the previous one, by an application of the transition rule (ActUnChg).

Please cite this article in press as: R. Lanotte, M. Merro, A semantic theory of the Internet of Things, Inf. Comput. (2018), https://doi.org/10.1016/j.ic.2018.01.001

- Let $M \rightarrowtriangle_\tau M'$ by an application of rule (parp):

$$\frac{\prod_i n_i[\mathtt{I}_i \bowtie P_i]^{\mu_i}_{h_i} \rightarrowtriangle_\tau \prod_i n_i[\mathtt{I}'_i \bowtie P'_i]^{\mu_i}_{h'_i}}{\prod_i n_i[\mathtt{I}_i \bowtie P_i \mid Q_i]^{\mu_i}_{h_i} \rightarrowtriangle_\tau \prod_i n_i[\mathtt{I}'_i \bowtie P'_i \mid Q_i]^{\mu_i}_{h'_i}}$$

  By inductive hypothesis we have: $\prod_i n_i[\mathtt{I}_i \bowtie P_i]^{\mu_i}_{h_i} \xrightarrow{\tau} \equiv \prod_i n_i[\mathtt{I}'_i \bowtie P'_i]^{\mu_i}_{h'_i}$. The τ-transition can be derived using different transition rules. Suppose that $\prod_i n_i[\mathtt{I}_i \bowtie P_i]^{\mu_i}_{h_i} \xrightarrow{\tau} \equiv \prod_i n_i[\mathtt{I}'_i \bowtie P'_i]^{\mu_i}_{h'_i}$ by an application of rule (SensRead) to node n_j, for some $j \in I$. Then, by using rule (ParP) to derive $P_j \mid Q_j \xrightarrow{s?v} P'_j \mid Q_j$, rule (SensRead) to derive $n_j[\mathtt{I}_j \bowtie P_j \mid Q_j]^{\mu_j}_{h_j} \xrightarrow{\tau} \equiv n_j[\mathtt{I}'_j \bowtie P'_j \mid Q_j]^{\mu_j}_{h'_j}$, and rule (ParN) to derive $\prod_i n_i[\mathtt{I}_i \bowtie P_i \mid Q_i]^{\mu_i}_{h_i} \xrightarrow{\tau} \equiv \prod_i n_i[\mathtt{I}'_i \bowtie P'_i \mid Q_i]^{\mu_i}_{h'_i}$, we get $M \xrightarrow{\tau} \equiv M'$.
  The cases when the τ-transition is derived by an application of the rules (ActUnChg), (Com) and (Pos) are similar.
- Let $M \rightarrowtriangle_\tau M'$ by an application of (loccom), with $M = n[\mathtt{I} \bowtie \lfloor \overline{c}\langle v\rangle.P \rfloor R \mid \lfloor c(x).Q \rfloor S]^{\mu}_{h}$ and $M' = n[\mathtt{I} \bowtie P \mid Q\{^v/_x\}]^{\mu}_{h}$, because $\mathrm{rng}(c) = -1$. Hence, for $\mathrm{rng}(c) = -1$ we can derive:

$$\frac{\dfrac{\lfloor \overline{c}\langle v\rangle.P \rfloor R \xrightarrow{\overline{c}v} P \qquad \lfloor c(v).Q \rfloor S \xrightarrow{cv} Q}{\lfloor \overline{c}\langle v\rangle.P \rfloor R \mid \lfloor c(x).Q \rfloor S \xrightarrow{\tau} P \mid Q\{^v/_x\}}}{n[\mathtt{I} \bowtie \lfloor \overline{c}\langle v\rangle.P \rfloor R \mid \lfloor c(x).Q \rfloor S]^{\mu}_{h} \xrightarrow{\tau} n[\mathtt{I} \bowtie P \mid Q\{^v/_x\}]^{\mu}_{h}}$$

  and $M \xrightarrow{\tau} \equiv M'$ is derived as required.
- Let $M \rightarrowtriangle_\tau M'$ by an application of (glbcom), with

$$M = n[\mathtt{I} \bowtie \lfloor \overline{c}\langle v\rangle.P \rfloor R]^{\mu}_{h} \mid m[\mathtt{I} \bowtie \lfloor c(x).Q \rfloor S]^{\mu'}_{k} \quad \text{and} \quad M' = n[\mathtt{I} \bowtie P]^{\mu}_{h} \mid m[\mathtt{I} \bowtie Q\{^v/_x\}]^{\mu'}_{k},$$

  because $\mathrm{d}(h,k) \leq \mathrm{rng}(c)$. Therefore the following derivation can be built up for $\mathrm{d}(h,k) \leq \mathrm{rng}(c)$:

$$\frac{\dfrac{\lfloor \overline{c}\langle v\rangle.P \rfloor R \xrightarrow{\overline{c}v} P}{n[\mathtt{I} \bowtie \lfloor \overline{c}\langle v\rangle.P \rfloor R]^{\mu}_{h} \xrightarrow{\overline{c}v@h} n[\mathtt{I} \bowtie P]^{\mu}_{h}} \qquad \frac{\lfloor c(v).Q \rfloor S \xrightarrow{cv} Q}{m[\mathtt{I} \bowtie \lfloor c(v).Q \rfloor S]^{\mu'}_{k} \xrightarrow{cv@k} m[\mathtt{I} \bowtie Q\{^v/_x\}]^{\mu'}_{k}}}{n[\mathtt{I} \bowtie \lfloor \overline{c}\langle v\rangle.P \rfloor R]^{\mu}_{h} \mid m[\mathtt{I} \bowtie \lfloor c(x).Q \rfloor S]^{\mu'}_{k} \xrightarrow{\tau} n[\mathtt{I} \bowtie P]^{\mu}_{h} \mid m[\mathtt{I} \bowtie Q\{^v/_x\}]^{\mu'}_{k}}$$

  and we get $M \xrightarrow{\tau} \equiv M'$.
- Let $M \rightarrowtriangle_\tau M'$ by an application of rule (res), with $M = (\boldsymbol{\nu}\tilde{g})M_1$ and $M' = (\boldsymbol{\nu}\tilde{g})M'_1$, because $M_1 \rightarrowtriangle_\tau M'_1$. By inductive hypothesis we have $M_1 \xrightarrow{\tau} \equiv M'_1$. Hence, by applying transition rules (Res), we can derive $M \xrightarrow{\tau} \equiv M'$.
- Let $M \rightarrowtriangle_\tau M'$ by an application of rule (struct) because $M \equiv N$, $N \rightarrowtriangle_\tau N'$ and $N' \equiv M'$. By inductive hypothesis we have $N \xrightarrow{\tau} \equiv N'$. Since $M \equiv N$ and $M' \equiv N'$, by an application of Lemma A.10 we obtain $M \xrightarrow{\tau} \equiv M'$.
- Let $M \rightarrowtriangle_\tau M'$ by an application of rule (parn), with $M = M_1 \mid N$ and $M' = M'_1 \mid N$, because $M_1 \rightarrowtriangle_\tau M'_1$. By inductive hypothesis $M_1 \rightarrowtriangle_\tau M'_1$ implies that $M_1 \xrightarrow{\tau} \equiv M'_1$. Hence, an application of the transition rule (ParN) concludes this case.

Let us prove the sub-result (3). The proof is by rule induction on why $M \xrightarrow{a} M'$.

- Let $M \xrightarrow{a} M'$ by an application of rule (ActChg), with $M = n[\mathtt{I} \bowtie P]^{\mu}_{h}$ and $M' = n[\mathtt{I}' \bowtie P']^{\mu}_{h}$, because $\mathtt{I}(a) = w \neq v$, $P \xrightarrow{a!v} P'$ and $\mathtt{I}' := \mathtt{I}[a \mapsto v]$. By Lemma A.8(3) there exist P_1, Q such that $P \equiv a!v.P_1 \mid Q$ and $P' \equiv P_1 \mid Q$. Then we can apply reduction rules (actchg) and (parp) to infer $M \rightarrowtriangle_a M'$.
- The cases when $M \xrightarrow{a} M'$ is derived by an application of either rule (ParN) or rule (Res) are analogous to the corresponding cases when $M \xrightarrow{\tau} M'$.

Let us prove the sub-result (4). The proof is by rule induction on why $M \rightarrowtriangle_a M'$.

- Let $M \rightarrowtriangle_a M'$ by an application of rule (actchg), with $M = n[\mathtt{I} \bowtie a!v.P]^{\mu}_{h}$ and $M' = n[\mathtt{I}' \bowtie P]^{\mu}_{h}$, because $\mathtt{I}(a) = w \neq v$ and $\mathtt{I}' = \mathtt{I}[a \mapsto v]$. By an application of rule (Actuator) we derive $a!v.P \xrightarrow{a!v} P$. The thesis follows by an application of rule (ActChg).
- The cases when $M \rightarrowtriangle_a M'$ is derived by an application of one of the rules among (parp), (parn), (res) or (struct) are analogous to the corresponding cases written for $M \rightarrowtriangle_\tau M'$.

Let us prove the sub-result (5). The proof is by rule induction on why $M \xrightarrow{\sigma} M'$.

- Let $M \xrightarrow{\sigma} M'$ by an application of rule (TimeZero). This case is immediate.
- Let $M \xrightarrow{\sigma} M'$ by an application of rule (TimeStat), with $M = n[\mathtt{I} \bowtie P]^{\mathrm{s}}_h$ and $M' = n[\mathtt{I} \bowtie P']^{\mathrm{s}}_h$, because $P \xrightarrow{\sigma} P'$ and $n[\mathtt{I} \bowtie P]^{\mathrm{s}}_h \xrightarrow{\tau}\!\!\!\!\!/\ \ $. Since $P \xrightarrow{\sigma} P'$, by Lemma A.8(1) we derive $P \equiv \prod_{i\in I} \lfloor \pi_i.P_i \rfloor Q_i \mid \prod_{j\in J} \sigma.P_j$ and $P' \equiv \prod_{i\in I} Q_i \mid \prod_{j\in J} P_j$ for some $I, J, \pi_i, P_i, Q_i, P_j$. By an application of the sub-result (2) above, from $n[\mathtt{I} \bowtie P]^{\mathrm{s}}_h \xrightarrow{\tau}\!\!\!\!\!/\ \ $ we derive $n[\mathtt{I} \bowtie P]^{\mathrm{s}}_h \not\rightarrowtriangle_\tau$. Then the thesis follows by applying the reduction rule (timestat).
- Let $M \xrightarrow{\sigma} M'$ by an application of rule (TimeMob). This case is similar to the previous one by applying the reduction rule (timemob) in place of (timestat).
- Let $M \xrightarrow{\sigma} M'$ by an application of rule (TimePar), with $M = M_1|M_2$ and $M' = M'_1|M'_2$, because $M_1 \xrightarrow{\sigma} M'_1$, $M_2 \xrightarrow{\sigma} M'_2$ and $M_1|M_2 \xrightarrow{\tau}\!\!\!\!\!/\ \ $. By inductive hypothesis we have $M_1 \rightarrowtriangle_\sigma M'_1$ and $M_2 \xrightarrow{\sigma} M'_2$. Moreover, by an application of the sub-result (2) above $M_1|M_2 \xrightarrow{\tau}\!\!\!\!\!/\ \ $ implies $M_1|M_2 \not\rightarrowtriangle_\tau$. Therefore we can apply the reduction rule (timepar) to get $M \rightarrowtriangle_\sigma M'$.
- Let $M \xrightarrow{\sigma} M'$ by an application of rule (Res). This case is similar to the corresponding one for $M \xrightarrow{\tau} M'$.

Let us prove the sub-result (6). The proof is by rule induction on why $M \rightarrowtriangle_\sigma M'$.

- Let $M \rightarrowtriangle_\sigma M'$ by an application of the reduction rule (timezero). This case is immediate.
- Let $M \rightarrowtriangle_\sigma M'$ by an application of rule (timestat):

$$\frac{n[\mathtt{I} \bowtie \prod_{i\in I} \lfloor \pi_i.P_i \rfloor Q_i \mid \prod_{j\in J} \sigma.P_j]^{\mathrm{s}}_h \not\rightarrowtriangle_\tau}{n[\mathtt{I} \bowtie \prod_{i\in I} \lfloor \pi_i.P_i \rfloor Q_i \mid \prod_{j\in J} \sigma.P_j]^{\mathrm{s}}_h \rightarrowtriangle_\sigma n[\mathtt{I} \bowtie \prod_{i\in I} Q_i \mid \prod_{j\in J} P_j]^{\mathrm{s}}_h}$$

  with $M = n[\mathtt{I} \bowtie \prod_{i\in I} \lfloor \pi_i.P_i \rfloor Q_i \mid \prod_{j\in J} \sigma.P_j]^{\mathrm{s}}_h$ and $M' = n[\mathtt{I} \bowtie \prod_{i\in I} Q_i \mid \prod_{j\in J} P_j]^{\mathrm{s}}_h$.
  By rule (Timeout) we derive $\lfloor \pi.P_i \rfloor Q_i \xrightarrow{\sigma} P_i$ and by rule (Delay) we derive $\sigma.P_j \xrightarrow{\sigma} P_j$. Now, we can repeatedly apply rule (TimeParP) to derive $\prod_{i\in I} \lfloor \pi_i.P_i \rfloor Q_i \mid \prod_{j\in J} \sigma.P_j \xrightarrow{\sigma} \prod_{i\in I} Q_i \mid \prod_{j\in J} P_j$. Indeed, by contradiction, if (TimeParP) would not be enabled, then rule (Com) would be enabled, and by applying rule (ParP), there would exist R such that $\prod_{i\in I} \lfloor \pi_i.P_i \rfloor Q_i \mid \prod_{j\in J} \sigma.P_j \xrightarrow{\tau} R$. Then, by applying rule (LocCom) and the sub-result (1) above, we would contradict the hypothesis $n[\mathtt{I} \bowtie \prod_{i\in I} \lfloor \pi_i.P_i \rfloor Q_i \mid \prod_{j\in J} \sigma.P_j]^{\mathrm{s}}_h \not\rightarrowtriangle_\tau$. Therefore, the thesis follows by applying the transition rule (TimeStat).
- Let $M \rightarrowtriangle_\sigma M'$ by an application of rule (timemob). This case is analogous to the previous one by applying the transition rule (TimeMob) in place of rule (TimeStat).
- Let $M \rightarrowtriangle_\sigma M'$ by an application of rule (timepar), with $M = M_1|M_2$ and $M' = M'_1|M'_2$, because $M_1 \rightarrowtriangle_\sigma M'_1$, $M_2 \rightarrowtriangle_\sigma M'_2$, and $M_1|M_2 \not\rightarrowtriangle_\tau$. By inductive hypothesis, $M_1 \rightarrowtriangle_\sigma M'_1$ implies $M_1 \xrightarrow{\sigma} \equiv M'_1$ and $M_2 \rightarrowtriangle_\sigma M'_2$ implies $M_2 \xrightarrow{\sigma} \equiv M'_2$. Finally, by an application of the sub-result (1) above $M_1|M_2 \not\rightarrowtriangle_\tau$ implies $M_1|M_2 \xrightarrow{\tau}\!\!\!\!\!/\ \ $. Therefore we can derive $M \xrightarrow{\sigma} \equiv M'$ by an application of the transition rule (TimePar).
- The cases when $M \rightarrowtriangle_\sigma M'$ is derived by an application of one of the rules (res) or (struct) are analogous to the corresponding cases written for $M \rightarrowtriangle_\tau M'$. □

A.4. Proofs of Section 5

Proof of Theorem 5.4. It remains to prove that the bisimilarity relation, $\approx$, is preserved by parallel composition and channel restriction.

Let us prove that $\approx$ is preserved by parallel composition. We show that the relation

$$\mathcal{R} = \{(M|O,\ N|O) : \text{ both } M|O \text{ and } N|O \text{ are well-formed and } M \approx N\}$$

is a bisimulation. We proceed by case analysis on why $M|O \xrightarrow{\alpha} \hat{M}$.

- Let $M|O \xrightarrow{\tau} \hat{M}$. We can distinguish two cases.
  - The transition is derived by applying rule (GlbCom), with $\hat{M} = M'|O'$, because $M \xrightarrow{\overline{c}v@h} M'$, $O \xrightarrow{cv@k} O'$, and $\mathrm{d}(h,k) \leq \mathrm{rng}(c)$. Since $M \xrightarrow{\overline{c}v@h} M'$ and $\mathrm{d}(h,k) \leq \mathrm{rng}(c)$, by an application of rule (SndObs) we derive $M \xrightarrow{\overline{c}v\triangleright k} M'$. As $M \approx N$, there are N_1, N_2 and N' such that $N \Longrightarrow N_1 \xrightarrow{\overline{c}v\triangleright k} N_2 \Longrightarrow N'$ with $M' \approx N'$. Thus, there exists a location h' such that $\mathrm{d}(h',k) \leq \mathrm{rng}(c)$ and $N_1 \xrightarrow{\overline{c}v@h'} N_2$. Therefore, by several applications of rule (ParN) and one application of rule (GlbCom) we can derive $N|O \Rightarrow \hat{N} = N'|O'$, with $(\hat{M}, \hat{N}) \in \mathcal{R}$. The symmetric case is analogous.
  - The transition is derived by applying rule (ParN) because $M \xrightarrow{\tau} M'$. As $M \approx N$ it follows that $N \Longrightarrow N'$ with $M' \approx N'$. By several applications of rule (ParN) it follows that $N|O \Rightarrow \hat{N} = N'|O'$, with $(\hat{M}, \hat{N}) \in \mathcal{R}$. The symmetric case is easier.
- Let $M|O \xrightarrow{\sigma} \hat{M} = M'|O'$. This is only possible by an application of rule (TimePar) because $M \xrightarrow{\sigma} M'$, $O \xrightarrow{\sigma} O'$ and $M|O \xrightarrow{\tau}\!\!\!\!\!/\ \ $. Since $M \approx N$ and $M \xrightarrow{\sigma} M'$ there are N_1, N_2 and N' such that $N \Rightarrow N_1 \xrightarrow{\sigma} N_2 \Rightarrow N'$, with $M' \approx N'$. By an

appropriate number of applications of rule (ParN) we have that $N|O \Rightarrow N_1|O$. Next step is to show that we can use rule (TimePar) to derive $N_1|O \xrightarrow{\sigma} N_2|O'$. For that we only need to prove that $N_1|O \not\xrightarrow{\tau}$. In fact, if $N_1|O \xrightarrow{\tau}$ then we would reach a contradiction. This because, $M \approx N$ and $N \Rightarrow N_1$ implies there is M_1 such that $M \Rightarrow M_1$ with $M_1 \approx N_1$. As $M \not\xrightarrow{\tau}$ it follows that $M = M_1 \approx N_1$. By Proposition 2.3, $N_1 \xrightarrow{\sigma} N_2$ and $O \xrightarrow{\sigma} O'$ imply $N_1 \not\xrightarrow{\tau}$ and $O \not\xrightarrow{\tau}$. Thus $N_1|O \xrightarrow{\tau}$ could be derived only by an application of rule (GlobCom) where N_1 interact with O via some channel c, with $\mathrm{rng}(c) \geq 0$. However, as $N_1 \approx M$ the network M could mimic the same interaction with O giving rise to a reduction of the form $M|O \Rightarrow \xrightarrow{\tau}$. This is in contradiction with the initial hypothesis that $M|O \not\xrightarrow{\tau}$. Thus, $N_1|O \not\xrightarrow{\tau}$ and by an application of rule (TimePar) we derive $N_1|O \xrightarrow{\sigma} N_2|O'$. By an appropriate number of applications of rule (ParN) we get $N_2|O' \Rightarrow N'|O'$. Thus, $N|O \stackrel{\sigma}{\Longrightarrow} \hat{N} = N'|O'$, with $(\hat{M}, \hat{N}) \in \mathcal{R}$.

- Let $M|O \xrightarrow{a} \hat{M}$. Then, we distinguish two cases.
  - Either $O \xrightarrow{a} O'$ and by an application of rule (ParN) we derive $M|O \xrightarrow{a} M|O'$. This case is easy.
  - Or $M \xrightarrow{a} M'$ and by an application of rule (ParN) we derive $M|O \xrightarrow{a} M'|O$. As $M \approx N$ there is N' such that $N \stackrel{a}{\Longrightarrow} N'$ and $M' \approx N'$. Thus, by several applications of rule (ParN) we derive $N|O \stackrel{a}{\Longrightarrow} \hat{N} = N'|O$, with $(\hat{M}, \hat{N}) \in \mathcal{R}$.
- Let $M|O \xrightarrow{\overline{c}v \triangleright k} \hat{M}$. By definition of rule (SndObs) this is only possible if $M|O \xrightarrow{\overline{c}v@h} \hat{M}$, with $\mathrm{d}(h,k) \leq \mathrm{rng}(c)$. Then, we distinguish two cases.
  - Either $O \xrightarrow{\overline{c}v@h} O'$ and $\hat{M} = M|O'$ by an application of rule (ParN). Then, by an application of the same rule we derive $N|O \xrightarrow{\overline{c}v@h} N|O'$. By an application of rule (SndObs) we get $N|O \xrightarrow{\overline{c}v \triangleright k} \hat{N} = N|O'$, with $(\hat{M}, \hat{N}) \in \mathcal{R}$.
  - Or $M \xrightarrow{\overline{c}v@h} M'$ and $\hat{M} = M'|O$ by an application of rule (ParN). By an application of rule (SndObs) we have $M \xrightarrow{\overline{c}v \triangleright k} M'$. As $M \approx N$ there is N' such that $N \stackrel{\overline{c}v \triangleright k}{\Longrightarrow} N'$, with $M' \approx N'$. As the transition $\xrightarrow{\overline{c}v \triangleright k}$ can only be derived by an application of rule (SndObs), it follows that $N \stackrel{\overline{c}v@h'}{\Longrightarrow} N'$, for some h' such that $\mathrm{d}(h',k) \leq \mathrm{rng}(c)$. By several applications of rule (ParN) it follows that $N|O \stackrel{\overline{c}v@h'}{\Longrightarrow} N'|O$. By an application of rule (SndObs) we finally obtain $N|O \stackrel{\overline{c}v \triangleright k}{\Longrightarrow} \hat{N} = N'|O$, with $(\hat{M}, \hat{N}) \in \mathcal{R}$.
- Let $M|O \xrightarrow{cv \triangleright k} \hat{M}$. This case is similar to the previous one.

Let us prove that $\approx$ is preserved by channel restriction. We show that the relation $\mathcal{R}$, defined as $\{\big((\boldsymbol{\nu} c)M,\ (\boldsymbol{\nu} c)N\big) : M \approx N\}$ is a bisimulation. We proceed by case analysis on why $(\boldsymbol{\nu} c)M \xrightarrow{\alpha} \hat{M}$.

- Let $(\boldsymbol{\nu} c)M \xrightarrow{\alpha} \hat{M}$, for $\alpha \in \{\tau, \sigma, a\}$. In this case, this transition has been derived by an application of rule (Res) because $M \xrightarrow{\alpha} M'$, with $\hat{M} = (\boldsymbol{\nu} c)M'$. As $M \approx N$ there is N' such that $N \stackrel{\alpha}{\Longrightarrow} N'$ and $M' \approx N'$. By several applications of rule (Res) we can derive $(\boldsymbol{\nu} c)N \stackrel{\alpha}{\Longrightarrow} \hat{N} = (\boldsymbol{\nu} c)N'$, with $(\hat{M}, \hat{N}) \in \mathcal{R}$.
- Let $(\boldsymbol{\nu} c)M \xrightarrow{\alpha} \hat{M}$, for $\alpha \in \{\overline{d}v \triangleright k, dv \triangleright k\}$, with $d \neq c$. This case is similar to the previous one except for the fact that we need to pass through the definitions of rules (SndObs) and (RcvObs) as the rule (Res) is only defined for intensional actions.
- Let $(\boldsymbol{\nu} c)M \xrightarrow{\alpha} \hat{M}$, for $\alpha \in \{\overline{c}v \triangleright k, cv \triangleright k\}$. This case is not admissible as rule (Res) blocks intensional actions of the form $\overline{c}v@h$ and $cv@h$. □

Next, we prove Lemma 5.7. For that we need a technical lemma.

Lemma A.11. *Let $O = n[\mathtt{I} \bowtie \mathsf{nil}]^{\mathrm{s}}_k$ for an arbitrary node name n, an arbitrary actuator a, and an arbitrary value v in the domain of a, such that $\mathtt{I}$ is only defined for a and $\mathtt{I}(a) = v$. If $M|O \cong N|O$ then $M \cong N$.*

Proof. We recall that we always work with well-formed systems. The proofs consists in showing that the relation

$$\mathcal{R} = \{(M, N) : M|O \cong N|O, \text{ for some } O \text{ defined as above}\}$$

is barb preserving, reduction closed and contextual. Since $\cong$ is the largest relation satisfying these properties, then $\mathcal{R} \subseteq \cong$ and therefore $M \cong N$. The scheme of the proof is very similar to that of the following proof. □

Proof of Lemma 5.7. Let $O = n[\mathtt{I} \bowtie a!v.\mathsf{nil}]^{\mathrm{s}}_k$, for an arbitrary node name n, an arbitrary actuator a, and arbitrary values v and w, in the domain of a, such that $\mathtt{I}$ is only defined for a and $\mathtt{I}(a) = w \neq v$. Let use define the relation

$$\mathcal{R} = \{(M, N) : M|O \cong N|O, \text{for some } O \text{ defined as above}\}.$$

We show that the relation $\mathcal{R} \cup \cong$ is barb preserving, reduction closed and contextual. Since $\cong$ is the largest relation satisfying these properties, then $\mathcal{R} \subseteq \cong$ and therefore $M \cong N$.

We recall that in this paper we only consider well-formed networks. So, in the definition of $\mathcal{R}$ we assume that all networks of the form $M|O$ and $N|O$ are well-formed. In particular, in order to decide whether $(M,N)\in\mathcal{R}$ it is enough to find an O of the indicated shape, which respects the requirements of $\mathcal{R}$, and which preserves well-formedness.

Let us prove that $\mathcal{R}\cup\cong$ is barb-preserving. We concentrate on the relation $\mathcal{R}$. As O has neither channels or sensors it is basically isolated from the rest of the world, except for signals emitted on the actuator a. So, it is very easy to see that $\mathcal{R}$ is barb preserving from $M|O\cong N|O$.

Let us prove that $\mathcal{R}\cup\cong$ is reduction closed. We focus on $\mathcal{R}$. Recall that $\rightarrow \stackrel{\text{def}}{=} \rightarrow_\tau\cup\rightarrow_\sigma$. Let $(M,N)\in\mathcal{R}$ and $M\rightarrow_\tau M'$, for some M'. We have to show that $N\rightarrow^* N'$, for some N' such that $(M',N')\in\mathcal{R}\cup\cong$. Let us fix an O which respects the requirements of $\mathcal{R}$. By an application of rule (parn) we infer $M|O\rightarrow_\tau M'|O$. As $M|O\cong N|O$ there is $\overline{N}$ such that $N|O\rightarrow^*\overline{N}$ and $M'|O\cong\overline{N}$. Since O cannot communicate and since the only enabled reduction for O is $\rightarrow_a$, none of the reductions in the reduction sequence $N|O\rightarrow^*\overline{N}$ involves O and none of these reductions is a timed one. Therefore, $\overline{N}=N'|O$, $N\rightarrow_\tau^* N'$, and $M'|O\cong N'|O$. This implies $(M',N')\in\mathcal{R}$.

Let $(M,N)\in\mathcal{R}$ and $M\rightarrow_b M'$, for some M'. As both networks $M|O$ and $N|O$ are well-formed, the actuator a cannot appear neither in M or in N. Thus, $a\neq b$. Starting from $M|O\cong N|O$ we reason as in the previous case.

Let $(M,N)\in\mathcal{R}$ and $M\rightarrow_\sigma M'$, for some M'. We have to show that $N\rightarrow^* N''$, for some N'' such that $(M',N'')\in\mathcal{R}\cup\cong$. By definition of $\mathcal{R}$ we have $M|O\cong N|O$. Let $M|O\rightarrow_a M|n[\mathtt{I}[a\mapsto v]\bowtie\mathsf{nil}]_k^{\mathsf{s}}$, by an application of rules (actchg) and (parn). As $\cong$ is reduction closed it follows that there is $\overline{N}$ such that $N|O\rightarrow^*\rightarrow_a\rightarrow^*\overline{N}$, with $M|n[\mathtt{I}[a\mapsto v]\bowtie\mathsf{nil}]_k^{\mathsf{s}}\cong\overline{N}$. Due to the structure of O the last reduction sequence can be decomposed as follows: $N|O\rightarrow^*\rightarrow_a\rightarrow^*\overline{N}=N'|n[\mathtt{I}[a\mapsto v]\bowtie\mathsf{nil}]_k^{\mathsf{s}}$, for some N' such that $N\rightarrow^* N'$. Thus, for $O'=n[\mathtt{I}[a\mapsto v]\bowtie\mathsf{nil}]_k^{\mathsf{s}}$, we have $M|O'\cong N'|O'$. Since $M\rightarrow_\sigma M'$, by Proposition 2.3, there is no M'' such that $M\rightarrow_\tau M''$. More generally, by looking at the definition of O' it is easy to see that there is no U such that $M|O'\rightarrow_\tau U$. Thus, by an application of rules (timestat) and (timepar) we can infer $M|O'\rightarrow_\sigma M'|O'$. As $M|O'\cong N'|O'$, by Proposition 2.14 there is $\hat{N}$ such that $N'|O'\rightarrow_\tau^*\rightarrow_\sigma\rightarrow_\tau^*\hat{N}$ and $M'|O'\cong\hat{N}$. By looking at the definition of O' the only possibility is that $\hat{N}=N''|O'$, with $N'\rightarrow_\tau^*\rightarrow_\sigma\rightarrow_\tau^* N''$ and $M'|O'\cong N''|O'$. By Lemma A.11 this implies $M'\cong N''$. Recapitulating we have that for $M\rightarrow_\sigma M'$ there is N'' such that $N\rightarrow^* N''$, with $(M',N'')\in\mathcal{R}\cup\cong$.

Let us prove that $\mathcal{R}\cup\cong$ is contextual. Again, it is enough to focus on $\mathcal{R}$. Let us consider the three different network contexts:

- Let $(M,N)\in\mathcal{R}$. Let O' be an arbitrary network such that both $M|O'$ and $N|O'$ are well formed. We want to show that $(M|O',\ N|O')\in\mathcal{R}$. As $(M,N)\in\mathcal{R}$, we can always find an $O=n[\mathtt{I}\bowtie a!v.\mathsf{nil}]_k^{\mathsf{s}}$ which respects the requirements of $\mathcal{R}$ such that $M|O\cong N|O$ and both networks $M|O|O'$ and $N|O|O'$ are well-formed. As $\cong$ is contextual and structural congruence is a monoid with respect to parallel composition, it follows that $(M|O')|O\equiv(M|O)|O'\cong(N|O)|O'\equiv(N|O')|O$. As $\equiv\subset\cong$ and $\cong$ is trivially transitive, this is enough to derive that $(M|O',\ N|O')\in\mathcal{R}$.
- Let $(M,N)\in\mathcal{R}$. Let c be an arbitrary channel name. Let $O=n[\mathtt{I}\bowtie a!v.\mathsf{nil}]_k^{\mathsf{s}}$ which respects the requirements of $\mathcal{R}$. As $\cong$ is contextual if follows that $(\boldsymbol{\nu}c)(M|O)\cong(\boldsymbol{\nu}c)(N|O)$. Since O does not contain channels it holds that $((\boldsymbol{\nu}c)M)|O\equiv(\boldsymbol{\nu}c)(M|O)\cong(\boldsymbol{\nu}c)(N|O)\equiv((\boldsymbol{\nu}c)N)|O$. As $\equiv\subset\cong$ and $\cong$ is trivially transitive, this is enough to derive that $((\boldsymbol{\nu}c)M,(\boldsymbol{\nu}c)N)\in\mathcal{R}$.
- Let $(M,N)\in\mathcal{R}$. Let $O=n[\mathtt{I}\bowtie a!v.\mathsf{nil}]_k^{\mathsf{s}}$ which respects the requirements of $\mathcal{R}$. Since O does not contain sensors, by Definition 2.9 we have: $M[s@h\mapsto v]|O=(M|O)[s@h\mapsto v]\cong(N|O)[s@h\mapsto v]=N[s@h\mapsto v]|O$. This is enough to derive that $(M[s@h\mapsto v],N[s@h\mapsto v])\in\mathcal{R}$. □

Proof of Theorem 5.11. For each law we exhibit the proper bisimulation. It is easy to see that for the first four laws the left-hand-side system evolves into the right-hand-side by performing a τ-actions. So, in order to prove these laws it is enough to show that the two terms under considerations are bisimilar. Let us proceed case by case.

1. Let use define the relation

$$\mathcal{R}=\left\{\left(n[\mathtt{I}\bowtie a!v.P|R]_h^\mu,\ n[\mathtt{I}\bowtie P|R]_h^\mu\right)\mid\mathtt{I}(a)=v\text{ and }a\text{ does not occur in }R\right\}\cup Id$$

   where Id is the identity relation. It suffices to prove the symmetric closure of $\mathcal{R}$ is a bisimulation.
2. Let us define the relation $\mathcal{R}=\{\left(n[\mathtt{I}\bowtie@(x).P|R]_h^\mu,\ n[\mathtt{I}\bowtie P\{{}^h/_x\}|R]_h^\mu\right)\}\cup Id$, where Id is the identity relation. We show that the symmetric closure of $\mathcal{R}$ is a bisimulation. The proof is similar to that of Law 1 where $n[\mathtt{I}\bowtie@(x).P|R]_h^\mu\xrightarrow{\tau}n[\mathtt{I}\bowtie P\{{}^h/_x\}|R]_h^\mu$.
3. Let us define the relation $\mathcal{R}=\left\{\left(n[\mathtt{I}\bowtie\lfloor\overline{c}\langle v\rangle.P\rfloor S|\lfloor c(x).Q\rfloor T|R]_h^\mu,\ n[\mathtt{I}\bowtie P|Q\{{}^v/_x\}|R]_h^\mu\right)\right\}\cup Id$, such that c is not in R and $\mathrm{rng}(c)=-1$. It suffices to show that the symmetric closure of $\mathcal{R}$ is a bisimulation.
4. The proof of Law 4 is similar to that of Law 3.
5. Let us define the relation $\mathcal{R}=\{\left(n[\mathtt{I}\bowtie P]_h^\mu,n[\mathtt{I}\bowtie\mathsf{nil}]_h^\mu\right)\}$, where P does not contains terms of the form $\lfloor\pi.P_1\rfloor P_2$ or $a!v.P_1$, for any a. It suffices to prove that the symmetric closure of $\mathcal{R}$ is a bisimulation.
6. Let us consider the relation $\mathcal{R}=\left\{\left(n[\mathtt{I}\bowtie\mathsf{nil}]_h^\mu,\mathbf{0}\right)\mid\mathtt{I}(a)\text{ is undefined for any actuator }a\right\}$. It suffices to prove that the symmetric closure of $\mathcal{R}$ is a bisimulation.

7. Let us define the relation $\mathcal{R} = \{(n[\emptyset \bowtie P]^{\mathrm{m}}_h, m[\emptyset \bowtie P]^{\mathrm{s}}_k)\}$, such that P does not contain terms $@(x).Q$ and for any channel c in P either $\mathrm{rng}(c) = \infty$ or $\mathrm{rng}(c) = -1$. It suffices to prove that the symmetric closure of $\mathcal{R}$ is a bisimulation. □

Proof of Lemma 5.12. Let use define the relation $\mathcal{R}$ as follows:

$$\left\{\left((\nu\tilde{c})(n[\mathtt{I} \bowtie P_1|R]^{\mu}_h|O_1), (\nu\tilde{d})(n[\mathtt{I} \bowtie P_2|R]^{\mu}_k|O_2)\right) : (\nu\tilde{c})(n[\mathtt{I} \bowtie P_1]^{\mu}_h|O_1) \approx (\nu\tilde{d})(n[\mathtt{I} \bowtie P_2]^{\mu}_k|O_2)\right\}$$

where process R can only (i) read the sensors of $\mathtt{I}$; (ii) transmit along some fresh Internet channel; (iii) let time passes. We prove that the symmetric closure of the relation $\mathcal{R}$ is a bisimulation. Let $(M, N) \in \mathcal{R}$, we proceed by case analysis on why $M \xrightarrow{\alpha} M'$.

- Let $M = (\nu\tilde{c})(n[\mathtt{I} \bowtie P_1|R]^{\mu}_h|O_1) \xrightarrow{\alpha} (\nu\tilde{c})(n[\mathtt{I}' \bowtie P'_1|R]^{\mu}_{h'}|O'_1) = M'$, with $\alpha \neq \sigma$, be a transitions which does not involve (and affect) R at all. This means that $(\nu\tilde{c})(n[\mathtt{I} \bowtie P_1]^{\mu}_h|O_1) \xrightarrow{\alpha} (\nu\tilde{c})(n[\mathtt{I}' \bowtie P'_1]^{\mu}_{h'}|O'_1)$. By hypothesis there are $\mathtt{I}''$, P'_2, O'_2 and k' such that $(\nu\tilde{c})(n[\mathtt{I} \bowtie P_2]^{\mu}_k|O_2) \stackrel{\alpha}{\Longrightarrow} (\nu\tilde{c})(n[\mathtt{I}'' \bowtie P'_2]^{\mu}_{k'}|O'_2)$ and $(\nu\tilde{c})(n[\mathtt{I}' \bowtie P'_1]^{\mu}_{h'}|O'_1) \approx (\nu\tilde{d})(n[\mathtt{I}'' \bowtie P'_2]^{\mu}_{k'}|O'_2)$. By Theorem 5.6 and Proposition 2.16 it follows that $\mathtt{I}' = \mathtt{I}''$. Furthermore as $\alpha \neq \sigma$ we have $h = h'$ and $k = k'$. Thus, $N = (\nu\tilde{c})(n[\mathtt{I} \bowtie P_2|R]^{\mu}_k|O_2) \stackrel{\hat{\alpha}}{\Longrightarrow} (\nu\tilde{c})(n[\mathtt{I}' \bowtie P'_2|R]^{\mu}_{k'}|O'_2) = N'$, with $(M', N') \in \mathcal{R}$.
- Let $M = (\nu\tilde{c})(n[\mathtt{I} \bowtie P_1|R]^{\mu}_h|O_1) \xrightarrow{\sigma} (\nu\tilde{c})(n[\mathtt{I} \bowtie P'_1|R']^{\mu}_{h'}|O'_1) = M'$. We know that timed actions do not change the physical interface $\mathtt{I}$. This implies that: (i) $R \xrightarrow{\sigma} R'$; (ii) $(\nu\tilde{c})(n[\mathtt{I} \bowtie P_1|R]^{\mu}_h|O_1) \stackrel{\tau}{\not\rightarrow}$; (iii) $(\nu\tilde{c})(n[\mathtt{I} \bowtie P_1]^{\mu}_h|O_1) \xrightarrow{\sigma} (\nu\tilde{c})(n[\mathtt{I} \bowtie P'_1]^{\mu}_{h'}|O'_1)$. In particular, the second item means that R does not have any interaction with the network. It even does not read some sensor of $\mathtt{I}$. By hypothesis we have that $(\nu\tilde{c})(n[\mathtt{I} \bowtie P_2]^{\mu}_k|O_2) \Rightarrow \xrightarrow{\sigma} \Rightarrow (\nu\tilde{c})(n[\mathtt{I}' \bowtie P'_2]^{\mu}_{k'}|O'_2)$ with $(\nu\tilde{c})(n[\mathtt{I} \bowtie P'_1]^{\mu}_{h'}|O'_1) \approx (\nu\tilde{d})(n[\mathtt{I}' \bowtie P'_2]^{\mu}_{k'}|O'_2)$. By Theorem 5.6 and Proposition 2.16 we know that it must be $\mathtt{I} = \mathtt{I}'$. As R cannot have any interaction with the rest of the network, apart from time synchronisation, it follows that $N = (\nu\tilde{c})(n[\mathtt{I} \bowtie P_2|R]^{\mu}_k|O_2) \Rightarrow \xrightarrow{\sigma} \Rightarrow (\nu\tilde{c})(n[\mathtt{I} \bowtie P'_2|R']^{\mu}_{k'}|O'_2) = N'$, with $(M', N') \in \mathcal{R}$.
- Let $M = (\nu\tilde{c})(n[\mathtt{I} \bowtie P_1|R]^{\mu}_h|O_1) \xrightarrow{\alpha} (\nu\tilde{c})(n[\mathtt{I} \bowtie P_1|R']^{\mu}_h|O'_1) = M'$, with $\alpha \neq \sigma$, be a transitions which is due to R. This can be a sensor reading or a transmission along some channel b, with $\mathrm{rng}(b) = \infty$. In that case, it is easy to see that, as $\mathrm{rng}(b) = \infty$, then $N = (\nu\tilde{c})(n[\mathtt{I} \bowtie P_2|R]^{\mu}_k|O_2) \xrightarrow{\alpha} (\nu\tilde{c})(n[\mathtt{I} \bowtie P_2|R']^{\mu}_k|O_2) = N'$, with $(M', N') \in \mathcal{R}$. □

Proof of Proposition 5.13. Let us introduce shorthands: $L \stackrel{\mathrm{def}}{=} \mathit{LightCtrl}$, $\overline{L} \stackrel{\mathrm{def}}{=} \overline{\mathit{LightCtrl}}$, $L_1 \stackrel{\mathrm{def}}{=} \mathit{LightMng}_1$, $L_2 \stackrel{\mathrm{def}}{=} \mathit{LightMng}_2$, and $\overline{C} \stackrel{\mathrm{def}}{=} \overline{\mathit{CLightMng}}$. Let us define the relation

$$\mathcal{R} \stackrel{\mathrm{def}}{=} \bigcup_{i=1}^{17} \left((\nu\tilde{c})M_i\,,\ (\nu\tilde{c})(\nu g)N_i\right)$$

where the pairs (M_i, N_i), for $1 \leq i \leq 17$, are listed below:

- $M_1 = n_P[\mathtt{I}_P \bowtie L]^{\mathrm{m}}_k \mid n_1[\mathtt{I}_1 \bowtie L_1]^{\mathrm{s}}_{loc1} \mid n_2[\mathtt{I}_2 \bowtie L_2]^{\mathrm{s}}_{loc4}$
  $N_1 = n_P[\mathtt{I}_P \bowtie \sigma.\overline{L}]^{\mathrm{m}}_k \mid n_1[\mathtt{I}_1 \bowtie L_1]^{\mathrm{s}}_{loc1} \mid n_2[\mathtt{I}_2 \bowtie L_2]^{\mathrm{s}}_{loc4} \mid n_C[\emptyset \bowtie \sigma.\overline{C}]^{\mathrm{s}}_{loc3}$, with $k \notin \{loc1, loc2, loc3, loc4\}$, $\mathtt{I}_1(\mathit{light}_1) = \mathsf{off}$ and $\mathtt{I}_2(\mathit{light}_2) = \mathsf{off}$
- $M_2 = n_P[\mathtt{I}_P \bowtie \sigma.L]^{\mathrm{m}}_{loc1} \mid n_1[\mathtt{I}_1 \bowtie \mathit{light}_1!\mathsf{on}.\sigma.L_1]^{\mathrm{s}}_{loc1} \mid n_2[\mathtt{I}_2 \bowtie L_2]^{\mathrm{s}}_{loc4}$
  $N_2 = n_P[\mathtt{I}_P \bowtie \sigma.\overline{L}]^{\mathrm{m}}_{loc1} \mid n_1[\mathtt{I}_1 \bowtie \mathit{light}_1!\mathsf{on}.\sigma.L_1]^{\mathrm{s}}_{loc1} \mid n_2[\mathtt{I}_2 \bowtie L_2]^{\mathrm{s}}_{loc4} \mid n_C[\emptyset \bowtie \sigma.\overline{C}]^{\mathrm{s}}_{loc3}$
- $M_3 = n_P[\mathtt{I}_P \bowtie \sigma.L]^{\mathrm{m}}_{loc1} \mid n_1[\mathtt{I}'_1 \bowtie \sigma.L_1]^{\mathrm{s}}_{loc1} \mid n_2[\mathtt{I}_2 \bowtie L_2]^{\mathrm{s}}_{loc4}$
  $N_3 = n_P[\mathtt{I}_P \bowtie \sigma.\overline{L}]^{\mathrm{m}}_{loc1} \mid n_1[\mathtt{I}'_1 \bowtie \sigma.L_1]^{\mathrm{s}}_{loc1} \mid n_2[\mathtt{I}_2 \bowtie L_2]^{\mathrm{s}}_{loc4} \mid n_C[\emptyset \bowtie \sigma.\overline{C}]^{\mathrm{s}}_{loc3}$, with $\mathtt{I}'_1(\mathit{light}_1) = \mathsf{on}$.
- $M_4 = n_P[\mathtt{I}_P \bowtie L]^{\mathrm{m}}_k \mid n_1[\mathtt{I}'_1 \bowtie L_1]^{\mathrm{s}}_{loc1} \mid n_2[\mathtt{I}_2 \bowtie L_2]^{\mathrm{s}}_{loc4}$
  $N_4 = n_P[\mathtt{I}_P \bowtie \sigma.\overline{L}]^{\mathrm{m}}_k \mid n_1[\mathtt{I}'_1 \bowtie L_1]^{\mathrm{s}}_{loc1} \mid n_2[\mathtt{I}_2 \bowtie L_2]^{\mathrm{s}}_{loc4} \mid n_C[\emptyset \bowtie \sigma.\overline{C}]^{\mathrm{s}}_{loc3}$, with $k \notin \{loc1, loc2, loc3, loc4\}$
- $M_5 = n_P[\mathtt{I}_P \bowtie L]^{\mathrm{m}}_k \mid n_1[\mathtt{I}'_1 \bowtie \mathit{light}_1!\mathsf{off}.L_1]^{\mathrm{s}}_{loc1} \mid n_2[\mathtt{I}_2 \bowtie L_2]^{\mathrm{s}}_{loc4}$
  $N_5 = n_P[\mathtt{I}_P \bowtie \sigma.\overline{L}]^{\mathrm{m}}_k \mid n_1[\mathtt{I}'_1 \bowtie \mathit{light}_1!\mathsf{off}.L_1]^{\mathrm{s}}_{loc1} \mid n_2[\mathtt{I}_2 \bowtie L_2]^{\mathrm{s}}_{loc4} \mid n_C[\emptyset \bowtie \sigma.\overline{C}]^{\mathrm{s}}_{loc3}$, with $k \notin \{loc1, loc2, loc3, loc4\}$
- $M_6 = n_P[\mathtt{I}_P \bowtie L]^{\mathrm{m}}_{loc1} \mid n_1[\mathtt{I}'_1 \bowtie \mathit{light}_1!\mathsf{off}.L_1]^{\mathrm{s}}_{loc1} \mid n_2[\mathtt{I}_2 \bowtie L_2]^{\mathrm{s}}_{loc4}$
  $N_6 = n_P[\mathtt{I}_P \bowtie \sigma.\overline{L}]^{\mathrm{m}}_{loc1} \mid n_1[\mathtt{I}'_1 \bowtie \mathit{light}_1!\mathsf{off}.L_1]^{\mathrm{s}}_{loc1} \mid n_2[\mathtt{I}_2 \bowtie L_2]^{\mathrm{s}}_{loc4} \mid n_C[\emptyset \bowtie \lfloor\overline{c_1}\langle\mathsf{on}\rangle.\sigma.\overline{C}\rfloor CLM]^{\mathrm{s}}_{loc3}$
- $M_7 = n_P[\mathtt{I}_P \bowtie L]^{\mathrm{m}}_{loc2} \mid n_1[\mathtt{I}'_1 \bowtie L_1]^{\mathrm{s}}_{loc1} \mid n_2[\mathtt{I}_2 \bowtie L_2]^{\mathrm{s}}_{loc4}$
  $N_7 = n_P[\mathtt{I}_P \bowtie \sigma.\overline{L}]^{\mathrm{m}}_{loc2} \mid n_1[\mathtt{I}'_1 \bowtie L_1]^{\mathrm{s}}_{loc1} \mid n_2[\mathtt{I}_2 \bowtie L_2]^{\mathrm{s}}_{loc4} \mid n_C[\emptyset \bowtie \sigma.\overline{C}]^{\mathrm{s}}_{loc3}$
- $M_8 = n_P[\mathtt{I}_P \bowtie L]^{\mathrm{m}}_{loc3} \mid n_1[\mathtt{I}'_1 \bowtie \mathit{light}_1!\mathsf{off}.L_1]^{\mathrm{s}}_{loc1} \mid n_2[\mathtt{I}_2 \bowtie L_2]^{\mathrm{s}}_{loc4}$
  $N_8 = n_P[\mathtt{I}_P \bowtie \sigma.\overline{L}]^{\mathrm{m}}_{loc3} \mid n_1[\mathtt{I}'_1 \bowtie \mathit{light}_1!\mathsf{off}.L_1]^{\mathrm{s}}_{loc1} \mid n_2[\mathtt{I}_2 \bowtie L_2]^{\mathrm{s}}_{loc4} \mid n_C[\emptyset \bowtie \sigma.\overline{C}]^{\mathrm{s}}_{loc3}$
- $M_9 = n_P[\mathtt{I}_P \bowtie L]^{\mathrm{m}}_{loc2} \mid n_1[\mathtt{I}'_1 \bowtie \mathit{light}_1!\mathsf{off}.L_1]^{\mathrm{s}}_{loc1} \mid n_2[\mathtt{I}_2 \bowtie L_2]^{\mathrm{s}}_{loc4}$
  $N_9 = n_P[\mathtt{I}_P \bowtie \sigma.\overline{L}]^{\mathrm{m}}_{loc2} \mid n_1[\mathtt{I}'_1 \bowtie \mathit{light}_1!\mathsf{off}.L_1]^{\mathrm{s}}_{loc1} \mid n_2[\mathtt{I}_2 \bowtie L_2]^{\mathrm{s}}_{loc4} \mid n_C[\emptyset \bowtie \sigma.\overline{C}]^{\mathrm{s}}_{loc3}$

- $M_{10} = n_P[\mathtt{I}_P \bowtie L]^{\mathrm{m}}_{loc3} \mid n_1[\mathtt{I}_1 \bowtie L_1]^{\mathrm{s}}_{loc1} \mid n_2[\mathtt{I}_2 \bowtie L_2]^{\mathrm{s}}_{loc4}$
  $N_{10} = n_P[\mathtt{I}_P \bowtie \sigma.\overline{L}]^{\mathrm{m}}_{loc3} \mid n_1[\mathtt{I}_1 \bowtie L_1]^{\mathrm{s}}_{loc1} \mid n_2[\mathtt{I}_2 \bowtie L_2]^{\mathrm{s}}_{loc4} \mid n_C[\emptyset \bowtie \sigma.\overline{C}]^{\mathrm{s}}_{loc3}$
- $M_{11} = n_P[\mathtt{I}_P \bowtie L]^{\mathrm{m}}_{loc2} \mid n_1[\mathtt{I}_1 \bowtie L_1]^{\mathrm{s}}_{loc1} \mid n_2[\mathtt{I}_2 \bowtie L_2]^{\mathrm{s}}_{loc4}$
  $N_{11} = n_P[\mathtt{I}_P \bowtie \sigma.\overline{L}]^{\mathrm{m}}_{loc2} \mid n_1[\mathtt{I}_1 \bowtie L_1]^{\mathrm{s}}_{loc1} \mid n_2[\mathtt{I}_2 \bowtie L_2]^{\mathrm{s}}_{loc4} \mid n_C[\emptyset \bowtie \sigma.\overline{C}]^{\mathrm{s}}_{loc3}$
- $M_{12} = n_P[\mathtt{I}_P \bowtie \sigma.L]^{\mathrm{m}}_{loc4} \mid n_1[\mathtt{I}_1 \bowtie L_1]^{\mathrm{s}}_{loc1} \mid n_2[\mathtt{I}_2 \bowtie light_2!\mathsf{on}.\sigma.L_2]^{\mathrm{s}}_{loc4}$
  $N_{12} = n_P[\mathtt{I}_P \bowtie \sigma.\overline{L}]^{\mathrm{m}}_{loc4} \mid n_1[\mathtt{I}_1 \bowtie L_1]^{\mathrm{s}}_{loc1} \mid n_2[\mathtt{I}_2 \bowtie light_2!\mathsf{on}.\sigma.L_2]^{\mathrm{s}}_{loc4} \mid n_C[\emptyset \bowtie \sigma.\overline{C}]^{\mathrm{s}}_{loc3}$
- $M_{13} = n_P[\mathtt{I}_P \bowtie \sigma.L]^{\mathrm{m}}_{loc4} \mid n_1[\mathtt{I}_1 \bowtie L_1]^{\mathrm{s}}_{loc1} \mid n_2[\mathtt{I}'_2 \bowtie \sigma.L_2]^{\mathrm{s}}_{loc4}$
  $N_{13} = n_P[\mathtt{I}_P \bowtie \sigma.\overline{L}]^{\mathrm{m}}_{loc4} \mid n_1[\mathtt{I}_1 \bowtie L_1]^{\mathrm{s}}_{loc1} \mid n_2[\mathtt{I}'_2 \bowtie \sigma.L_2]^{\mathrm{s}}_{loc4} \mid n_C[\emptyset \bowtie \sigma.\overline{C}]^{\mathrm{s}}_{loc3}$, where $\mathtt{I}'_2(light_2) = \mathsf{on}$
- $M_{14} = n_P[\mathtt{I}_P \bowtie L]^{\mathrm{m}}_{loc3} \mid n_1[\mathtt{I}_1 \bowtie L_1]^{\mathrm{s}}_{loc1} \mid n_2[\mathtt{I}'_2 \bowtie L_2]^{\mathrm{s}}_{loc4}$
  $N_{14} = n_P[\mathtt{I}_P \bowtie \sigma.\overline{L}]^{\mathrm{m}}_{loc3} \mid n_1[\mathtt{I}_1 \bowtie L_1]^{\mathrm{s}}_{loc1} \mid n_2[\mathtt{I}'_2 \bowtie L_2]^{\mathrm{s}}_{loc4} \mid n_C[\emptyset \bowtie \sigma.\overline{C}]^{\mathrm{s}}_{loc3}$
- $M_{15} = n_P[\mathtt{I}_P \bowtie L]^{\mathrm{m}}_{loc4} \mid n_1[\mathtt{I}_1 \bowtie L_1]^{\mathrm{s}}_{loc1} \mid n_2[\mathtt{I}'_2 \bowtie light_2!\mathsf{off}.L_2]^{\mathrm{s}}_{loc4}$
  $N_{15} = n_P[\mathtt{I}_P \bowtie \sigma.\overline{L}]^{\mathrm{m}}_{loc4} \mid n_1[\mathtt{I}_1 \bowtie L_1]^{\mathrm{s}}_{loc1} \mid n_2[\mathtt{I}'_2 \bowtie light_2!\mathsf{off}.L_2]^{\mathrm{s}}_{loc4} \mid n_C[\emptyset \bowtie \lfloor\overline{c_2}\langle\mathsf{on}\rangle.\sigma.\overline{C}\rfloor\overline{C}]^{\mathrm{s}}_{loc3}$
- $M_{16} = n_P[\mathtt{I}_P \bowtie L]^{\mathrm{m}}_{loc2} \mid n_1[\mathtt{I}_1 \bowtie L_1]^{\mathrm{s}}_{loc1} \mid n_2[\mathtt{I}'_2 \bowtie light_2!\mathsf{off}.L_2]^{\mathrm{s}}_{loc4}$
  $N_{16} = n_P[\mathtt{I}_P \bowtie \sigma.\overline{L}]^{\mathrm{m}}_{loc2} \mid n_1[\mathtt{I}_1 \bowtie L_1]^{\mathrm{s}}_{loc1} \mid n_2[\mathtt{I}'_2 \bowtie light_2!\mathsf{off}.L_2]^{\mathrm{s}}_{loc4} \mid n_C[\emptyset \bowtie \sigma.\overline{C}]^{\mathrm{s}}_{loc3}$
- $M_{17} = n_P[\mathtt{I}_P \bowtie L]^{\mathrm{m}}_{loc3} \mid n_1[\mathtt{I}_1 \bowtie L_1]^{\mathrm{s}}_{loc1} \mid n_2[\mathtt{I}'_2 \bowtie light_2!\mathsf{off}.L_2]^{\mathrm{s}}_{loc4}$
  $N_{17} = n_P[\mathtt{I}_P \bowtie \sigma.\overline{L}]^{\mathrm{m}}_{loc3} \mid n_1[\mathtt{I}_1 \bowtie L_1]^{\mathrm{s}}_{loc1} \mid n_2[\mathtt{I}'_2 \bowtie light_2!\mathsf{off}.L_2]^{\mathrm{s}}_{loc4} \mid n_C[\emptyset \bowtie \sigma.\overline{C}]^{\mathrm{s}}_{loc3}$.

For each pair $\big((\boldsymbol{\nu}\tilde{c})M_i\,,\,(\boldsymbol{\nu}\tilde{c},g)N_i\big) \in \mathcal{R}$ we proceed by case analysis on why $(\boldsymbol{\nu}\tilde{c})M_i \xrightarrow{\alpha} \hat{M}$. Then, we do the same for $(\boldsymbol{\nu}\tilde{c},g)N_i \xrightarrow{\alpha} \hat{N}$. Before starting the case analysis we notice that in all pairs of $\mathcal{R}$ the physical interfaces of the corresponding nodes are the same. For that reason we can safely omit the extensional actions of the form $a@h!v$. Moreover, our processes never read sensors (we removed from the initial system both the process *BoilerCtrl* and the network BM). Thus, we can safely omit transitions labelled with actions of the form $s@h?v$ as well.

– Let us consider the pair $\big((\boldsymbol{\nu}\tilde{c})M_1\,,\,(\boldsymbol{\nu}\tilde{c})(\boldsymbol{\nu}g)N_1\big)$. We proceed by case analysis on why $(\boldsymbol{\nu}\tilde{c})M_1 \xrightarrow{\alpha} \hat{M}$.

- Let $(\boldsymbol{\nu}\tilde{c})M_1 \xrightarrow{\alpha} \hat{M}$, for $\alpha \neq \sigma$. This case is not admissible as the phone is too far to interact with some local light manager.
- Let $(\boldsymbol{\nu}\tilde{c})M_1 \xrightarrow{\sigma} (\boldsymbol{\nu}\tilde{c})M'_1$, with

  $$M'_1 \;=\; n_P[\mathtt{I}_P \bowtie L]^{\mathrm{m}}_{k} \;\big|\; n_1[\mathtt{I}_1 \bowtie light_1!\mathsf{off}.L_1]^{\mathrm{s}}_{loc1} \;\big|\; n_2[\mathtt{I}_2 \bowtie light_2!\mathsf{off}.L_2]^{\mathrm{s}}_{loc4}$$

  and $k \notin \{loc1, loc2, loc3, loc4\}$. This means that the phone did not get inside the smart home. For the sake of simplicity we will call k all locations outside the smart home. By two applications of Law 1 of Theorem 5.11 we have: $(\boldsymbol{\nu}\tilde{c})M'_1 \gtrsim (\boldsymbol{\nu}\tilde{c})\big(n_P[\mathtt{I}_P \bowtie L]^{\mathrm{m}}_{k} \mid n_1[\mathtt{I}_1 \bowtie L_1]^{\mathrm{s}}_{loc1} \mid n_2[\mathtt{I}_2 \bowtie L_2]^{\mathrm{s}}_{loc4}\big) = (\boldsymbol{\nu}\tilde{c})M_1$. Then, $(\boldsymbol{\nu}\tilde{c},g)N_1 \xrightarrow{\sigma} \Rightarrow (\boldsymbol{\nu}\tilde{c},g)N_1$, and $\big((\boldsymbol{\nu}\tilde{c})M_1\,,\,(\boldsymbol{\nu}\tilde{c},g)N_1\big) \in \mathcal{R}$.
- Let $(\boldsymbol{\nu}\tilde{c})M_1 \xrightarrow{\sigma} (\boldsymbol{\nu}\tilde{c})M'_1$, with

  $$M'_1 \;=\; n_P[\mathtt{I}_P \bowtie L]^{\mathrm{m}}_{loc1} \;\big|\; n_1[\mathtt{I}_1 \bowtie light_1!\mathsf{off}.L_1]^{\mathrm{s}}_{loc1} \;\big|\; n_2[\mathtt{I}_2 \bowtie light_2!\mathsf{off}.L_2]^{\mathrm{s}}_{loc4}\,.$$

  In this case the smartphone just entered the smart home from its entrance, located at $loc1$. By two applications of Law 1 and one application of Law 4 of Theorem 5.11 we have:

  $$(\boldsymbol{\nu}\tilde{c})M'_1 \gtrsim (\boldsymbol{\nu}\tilde{c})\big(n_P[\mathtt{I}_P \bowtie \sigma.L]^{\mathrm{m}}_{loc1} \mid n_1[\mathtt{I}_1 \bowtie light_1!\mathsf{on}.\sigma.L_1]^{\mathrm{s}}_{loc1} \mid n_2[\mathtt{I}_2 \bowtie L_2]^{\mathrm{s}}_{loc4}\big) = (\boldsymbol{\nu}\tilde{c})M_2\,.$$

  Then, there is N_2 such that $(\boldsymbol{\nu}\tilde{c},g)N_1 \xrightarrow{\sigma} \Rightarrow (\boldsymbol{\nu}\tilde{c},g)N_2$ with

  $$(\boldsymbol{\nu}\tilde{c},g)N_2 \;=\; (\boldsymbol{\nu}\tilde{c},g)\big(n_P[\mathtt{I}_P \bowtie \sigma.\overline{L}]^{\mathrm{m}}_{loc1} \;\big|\; n_1[\mathtt{I}_1 \bowtie light_1!\mathsf{on}.\sigma.L_1]^{\mathrm{s}}_{loc1} \;\big|\; n_2[\mathtt{I}_2 \bowtie L_2]^{\mathrm{s}}_{loc4} \;\big|\; n_C[\emptyset \bowtie \sigma.\overline{C}]^{\mathrm{s}}_{loc3}\big)$$

  and $\big((\boldsymbol{\nu}\tilde{c})M_2\,,\,(\boldsymbol{\nu}\tilde{c},g)N_2\big) \in \mathcal{R}$.

Now, we proceed by case analysis on why $(\boldsymbol{\nu}\tilde{c},g)N_1 \xrightarrow{\alpha} \hat{N}$.

- Let $(\boldsymbol{\nu}\tilde{c},g)N_1 \xrightarrow{\alpha} \hat{N}$, with $\alpha \neq \sigma$. This case is not admissible.
- Let $(\boldsymbol{\nu}\tilde{c},g)N_1 \xrightarrow{\sigma} (\boldsymbol{\nu}\tilde{c},g)N'_1$, where the phone didn't enter the house, as its location is different from $loc1$. This case is similar to the previous one.
- $(\boldsymbol{\nu}\tilde{c},g)N_1 \xrightarrow{\sigma} (\boldsymbol{\nu}\tilde{c},g)N'_1$, with

  $$N'_1 \;=\; n_P[\mathtt{I}_P \bowtie \overline{L}]^{\mathrm{m}}_{loc1} \;\big|\; n_1[\mathtt{I}_1 \bowtie light_1!\mathsf{off}.L_1]^{\mathrm{s}}_{loc1} \;\big|\; n_2[\mathtt{I}_2 \bowtie light_2!\mathsf{off}.L_2]^{\mathrm{s}}_{loc4} \;\big|\; n_C[\emptyset \bowtie \overline{C}]^{\mathrm{s}}_{loc3}\,.$$

Because the phone just moved to location $loc1$. By two applications of Law 1, one application of Law 2, and two applications of Law 4 of Theorem 5.11 we have:

$$(\nu\tilde{c}, g)N_1' \gtrsim (\nu\tilde{c}, g)\big(n_P[\mathtt{I}_P \bowtie \sigma.\overline{L}]^{\mathtt{m}}_{loc1} \mid n_1[\mathtt{I}_1 \bowtie light_1!\mathsf{on}.\sigma.L_1]^{\mathtt{s}}_{loc1} \mid n_2[\mathtt{I}_2 \bowtie L_2]^{\mathtt{s}}_{loc4} \mid n_C[\emptyset \bowtie \sigma.\overline{C}]^{\mathtt{s}}_{loc3}\big) = (\nu\tilde{c}, g)N_2\,.$$

Then, there is M_2 such that $(\nu\tilde{c})M_1 \xrightarrow{\sigma} \Rightarrow (\nu\tilde{c})M_2$, with

$$M_2 = n_P[\mathtt{I}_P \bowtie \sigma.L]^{\mathtt{m}}_{loc1} \mid n_1[\mathtt{I}_1 \bowtie light_1!\mathsf{on}.\sigma.L_1]^{\mathtt{s}}_{loc1} \mid n_2[\mathtt{I}_2 \bowtie L_2]^{\mathtt{s}}_{loc4}$$

and $\big((\nu\tilde{c})M_2\,,\ (\nu\tilde{c}, g)N_2\big) \in \mathcal{R}$.

– Let us consider the pair $\big((\nu\tilde{c})M_2\,,\ (\nu\tilde{c})(\nu g)N_2\big)$. The only possible transition in both networks is a strong transition $\xrightarrow{light_1}$ which leads to the pair $\big((\nu\tilde{c})M_3\,,\ (\nu\tilde{c})(\nu g)N_3\big) \in \mathcal{R}$.

– Let us consider the pair $\big((\nu\tilde{c})M_3\,,\ (\nu\tilde{c})(\nu g)N_3\big)$. We proceed by case analysis on why $(\nu\tilde{c})M_3 \xrightarrow{\alpha} \hat{M}$.

- Let $(\nu\tilde{c})M_3 \xrightarrow{\alpha} \hat{M}$, for $\alpha \neq \sigma$. This case is not admissible.
- Let $(\nu\tilde{c})M_3 \xrightarrow{\sigma} (\nu\tilde{c})M_3'$, where $M_3' = n_P[\mathtt{I}_P \bowtie L]^{\mathtt{m}}_{loc1} \mid n_1[\mathtt{I}_1' \bowtie L_1]^{\mathtt{s}}_{loc1} \mid n_2[\mathtt{I}_2 \bowtie light_2!\mathsf{off}.L_2]^{\mathtt{s}}_{loc4}$ because the phone remained at location $loc1$. By two applications of Law 1 and one application of Law 4 of Theorem 5.11 we get:

  $$(\nu\tilde{c})M_3' \gtrsim (\nu\tilde{c})\big(n_P[\mathtt{I}_P \bowtie \sigma.L]^{\mathtt{m}}_{loc1} \mid n_1[\mathtt{I}_1' \bowtie \sigma.L_1]^{\mathtt{s}}_{loc1} \mid n_2[\mathtt{I}_2 \bowtie L_2]^{\mathtt{s}}_{loc4}\big) = (\nu\tilde{c})M_3\,.$$

  Then, $(\nu\tilde{c}, g)N_3 \xrightarrow{\sigma} \Rightarrow (\nu\tilde{c}, g)N_3$, and obviously $\big((\nu\tilde{c})M_3\,,\ (\nu\tilde{c}, g)N_3\big) \in \mathcal{R}$.
- Let $(\nu\tilde{c})M_3 \xrightarrow{\sigma} (\nu\tilde{c})M_3'$, where $M_3' = n_P[\mathtt{I}_P \bowtie L]^{\mathtt{m}}_{k} \mid n_1[\mathtt{I}_1' \bowtie L_1]^{\mathtt{s}}_{loc1} \mid n_2[\mathtt{I}_2 \bowtie light_2!\mathsf{off}.L_2]^{\mathtt{s}}_{loc4}$, with $k \notin \{loc1, loc2, loc3, loc4\}$, i.e. the phone moved out of the house. By applying Law 1 of Theorem 5.11 we get

  $$(\nu\tilde{c})M_3' \gtrsim (\nu\tilde{c})n_P[\mathtt{I}_P \bowtie L]^{\mathtt{m}}_{k} \mid n_1[\mathtt{I}_1' \bowtie L_1]^{\mathtt{s}}_{loc1} \mid n_2[\mathtt{I}_2 \bowtie L_2]^{\mathtt{s}}_{loc4} = (\nu\tilde{c})M_4\,.$$

  Then, there is N_4 such that $(\nu\tilde{c}, g)N_3 \xrightarrow{\sigma} \Rightarrow (\nu\tilde{c}, g)N_4$ with:

  $$N_4 = n_P[\mathtt{I}_P \bowtie \sigma.\overline{L}]^{\mathtt{m}}_{k} \mid n_1[\mathtt{I}_1' \bowtie L_1]^{\mathtt{s}}_{loc1} \mid n_2[\mathtt{I}_2 \bowtie L_2]^{\mathtt{s}}_{loc4} \mid n_C[\emptyset \bowtie \sigma.\overline{C}]^{\mathtt{s}}_{loc3}$$

  and $\big((\nu\tilde{c})M_4\,,\ (\nu\tilde{c}, g)N_4\big) \in \mathcal{R}$.
- Let $(\nu\tilde{c})M_3 \xrightarrow{\sigma} (\nu\tilde{c})M_3'$, where $M_3' = n_P[\mathtt{I}_P \bowtie L]^{\mathtt{m}}_{loc2} \mid n_1[\mathtt{I}_1' \bowtie L_1]^{\mathtt{s}}_{loc1} \mid n_2[\mathtt{I}_2 \bowtie light_2!\mathsf{off}.L_2]^{\mathtt{s}}_{loc4}$, because the phone moved from $loc1$ to $loc2$. In this case, by applying Law 1 of Theorem 5.11 we have:

  $$(\nu\tilde{c})M_3' \gtrsim (\nu\tilde{c})\big(n_P[\mathtt{I}_P \bowtie L]^{\mathtt{m}}_{loc2} \mid n_1[\mathtt{I}_1' \bowtie L_1]^{\mathtt{s}}_{loc1} \mid n_2[\mathtt{I}_2 \bowtie L_2]^{\mathtt{s}}_{loc4}\big) = (\nu\tilde{c})M_7\,.$$

  Then, we have that $(\nu\tilde{c}, g)N_3 \xrightarrow{\sigma} \Rightarrow (\nu\tilde{c}, g)N_7$, where

  $$N_7 = n_P[\mathtt{I}_P \bowtie \sigma.\overline{L}]^{\mathtt{m}}_{loc2} \mid n_1[\mathtt{I}_1' \bowtie L_1]^{\mathtt{s}}_{loc1} \mid n_2[\mathtt{I}_2 \bowtie L_2]^{\mathtt{s}}_{loc4} \mid n_C[\emptyset \bowtie \sigma.\overline{C}]^{\mathtt{s}}_{loc3}$$

  and $\big((\nu\tilde{c})M_7\,,\ (\nu\tilde{c}, g)N_7\big) \in \mathcal{R}$.

The case analysis when $(\nu\tilde{c}, g)N_3 \xrightarrow{\alpha} \hat{N}$ is similar.
The remaining cases, dealing with the pairs $\big((\nu\tilde{c})M_i\,,\ (\nu\tilde{c})(\nu g)N_i\big)$, for $4 \leq i \leq 17$, work in a similar manner. □

References

[1] R. Lanotte, M. Merro, A semantic theory of the Internet of Things (extended abstract), in: COORDINATION, in: Lect. Notes Comput. Sci., vol. 9686, Springer, 2016, pp. 157–174.
[2] J. Gubbi, R. Buyya, S. Marusic, M. Palaniswami, Internet of Things (IoT): a vision, architectural elements, and future directions, Future Gener. Comput. Syst. 29 (7) (2013) 1645–1660, https://doi.org/10.1016/j.future.2013.01.010.
[3] L. Atzori, A. Iera, G. Morabito, The Internet of Things: a survey, Comput. Netw. 54 (15) (2010) 2787–2805, https://doi.org/10.1016/j.comnet.2010.05.010.
[4] D. Miorandi, S. Sicari, F. De Pellegrini, I. Chlamtac, Internet of Things: vision, applications and research challenges, Ad Hoc Netw. 10 (7) (2012) 1497–1516, https://doi.org/10.1016/j.adhoc.2012.02.016.
[5] G. Roussos, V. Kostakos, RFID in pervasive computing: state-of-the-art and outlook, Pervasive Mob. Comput. 5 (1) (2009) 110–131, https://doi.org/10.1016/j.pmcj.2008.11.004.
[6] M.P. Papazoglou, W. van den Heuvel, Service oriented architectures: approaches, technologies and research issues, VLDB J. 16 (3) (2007) 389–415, https://doi.org/10.1007/s00778-007-0044-3.
[7] S. De, T. Elsaleh, P. Barnaghi, S. Meissner, An Internet of Things platform for real-world and digital objects, Scalable Comput. Pract. Exp. 13 (1).
[8] I. Lanese, L. Bedogni, M. Di Felice, Internet of Things: a process calculus approach, in: ACM SAC, ACM, 2013, pp. 1339–1346.
[9] K. Honda, N. Yoshida, On reduction-based process semantics, Theor. Comput. Sci. 151 (2) (1995) 437–486, https://doi.org/10.1016/0304-3975(95)00074-7.

Please cite this article in press as: R. Lanotte, M. Merro, A semantic theory of the Internet of Things, Inf. Comput. (2018), https://doi.org/10.1016/j.ic.2018.01.001

[10] D. Sangiorgi, D. Walker, The π-Calculus: a Theory of Mobile Processes, Cambridge University Press, 2001.
[11] A. van der Schaft, J. Schumacher, An Introduction to Hybrid Dynamical Systems, Lect. Notes Control Inf. Sci., vol. 251, Springer, 2000.
[12] G. Plotkin, A Structural Approach to Operational Semantics, Report DAIMI FN-19, Aarhus University, 1981.
[13] M. Hennessy, T. Regan, A process algebra for timed systems, Inf. Comput. 117 (2) (1995) 221–239, https://doi.org/10.1006/inco.1995.1041.
[14] R. Milner, D. Sangiorgi, Barbed bisimulation, in: ICALP, in: Lect. Notes Comput. Sci., vol. 623, Springer, 1992, pp. 685–695.
[15] A. Cerone, M. Hennessy, M. Merro, Modelling MAC-layer communications in wireless systems, Log. Methods Comput. Sci. 11 (1:18) (2015), https://doi.org/10.2168/LMCS-11(1:18)2015.
[16] B. Sundararaman, U. Buy, A. Kshemkalyani, Clock synchronization for wireless sensor networks: a survey, Ad Hoc Netw. 3 (3) (2005) 281–323, https://doi.org/10.1016/j.adhoc.2005.01.002.
[17] G. Berry, G. Gonthier, The Esterel synchronous programming language: design, semantics, implementation, Sci. Comput. Program. 19 (2) (1992) 87–152, https://doi.org/10.1016/0167-6423(92)90005-V.
[18] R. Amadio, A synchronous pi-calculus, Inf. Comput. 205 (9) (2007) 1470–1490, https://doi.org/10.1016/j.ic.2007.02.002.
[19] F. Boussinot, R. de Simone, The SL synchronous language, IEEE Trans. Softw. Eng. 22 (4) (1996) 256–266, https://doi.org/10.1109/32.491649.
[20] L. Cardelli, A. Gordon, Mobile ambients, Theor. Comput. Sci. 240 (1) (2000) 177–213, https://doi.org/10.1016/S0304-3975(99)00231-5.
[21] G. Wang, G. Cao, T. La Porta, Movement-assisted sensor deployment, IEEE Trans. Mob. Comput. 5 (6) (2006) 640–652, https://doi.org/10.1109/TMC.2006.80.
[22] M. Merro, F. Zappa Nardelli, Behavioral theory for mobile ambients, J. ACM 52 (6) (2005) 961–1023, https://doi.org/10.1145/1101821.1101825.
[23] S. Arun-Kumar, M. Hennessy, An efficiency preorder for processes, Acta Inform. 29 (8) (1992) 737–760, https://doi.org/10.1007/BF01191894.
[24] C. Bodei, P. Degano, G. Ferrari, L. Galletta, Where do your iot ingredients come from?, in: COORDINATION, in: Lect. Notes Comput. Sci., vol. 9686, Springer, 2016, pp. 35–50.
[25] C. Bodei, P. Degano, G. Ferrari, L. Galletta, Tracing where IoT data are collected and aggregated, Log. Methods Comput. Sci. 13 (3) (2017) 1–38, https://doi.org/10.23638/LMCS-13(3:5)2017.
[26] D. Gelernter, Generative communication in Linda, ACM Trans. Program. Lang. Syst. 7 (1) (1985) 80–112, https://doi.org/10.1145/2363.2433.
[27] R. Lanotte, M. Merro, A calculus of cyber-physical systems, in: LATA, vol. 10168, Springer, 2017, pp. 115–127.
[28] R. Lanotte, M. Merro, S. Tini, A probabilistic calculus of cyber-physical systems, CoRR abs/1707.02279.
[29] R. Lanotte, M. Merro, R. Muradore, L. Viganò, A formal approach to cyber-physical attacks, in: IEEE CSF, IEEE Computer Society, 2017, pp. 436–450.
[30] I. Lanese, D. Sangiorgi, An operational semantics for a calculus for wireless systems, Theor. Comput. Sci. 411 (2010) 1928–1948, https://doi.org/10.1016/j.tcs.2010.01.023.
[31] S. Nanz, C. Hankin, A framework for security analysis of mobile wireless networks, Theor. Comput. Sci. 367 (1–2) (2006) 203–227, https://doi.org/10.1016/j.tcs.2006.08.036.
[32] M. Merro, An observational theory for mobile ad hoc networks (full paper), Inf. Comput. 207 (2) (2009) 194–208, https://doi.org/10.1016/j.ic.2007.11.010.
[33] J. Godskesen, A calculus for mobile ad hoc networks, in: COORDINATION, in: Lect. Notes Comput. Sci., vol. 4467, Springer, 2007, pp. 132–150.
[34] F. Ghassemi, W. Fokkink, A. Movaghar, Verification of mobile ad hoc networks: an algebraic approach, Theor. Comput. Sci. 412 (28) (2011) 3262–3282, https://doi.org/10.1016/j.tcs.2011.03.017.
[35] M. Merro, F. Ballardin, E. Sibilio, A timed calculus for wireless systems, Theor. Comput. Sci. 412 (47) (2011) 6585–6611, https://doi.org/10.1016/j.tcs.2011.07.016.
[36] M. Merro, E. Sibilio, A calculus of trustworthy ad hoc networks, Form. Asp. Comput. 25 (5) (2013) 801–832, https://doi.org/10.1007/s00165-011-0210-7.
[37] R. Lanotte, M. Merro, Semantic analysis of gossip protocols for wireless sensor networks, in: CONCUR, in: Lect. Notes Comput. Sci., vol. 6901, Springer, 2011, pp. 156–170.
[38] A. Singh, C. Ramakrishnan, S. Smolka, A process calculus for mobile ad hoc networks, Sci. Comput. Program. 75 (6) (2010) 440–469, https://doi.org/10.1016/j.scico.2009.07.008.
[39] A. Fehnker, R. van Glabbeek, P. Höfner, A. McIver, M. Portmann, W. Tan, A process algebra for wireless mesh networks, in: ESOP, in: Lect. Notes Comput. Sci., vol. 7211, Springer, 2012, pp. 295–315.
[40] J. Borgström, S. Huang, M. Johansson, P. Raabjerg, B. Victor, J. Pohjola, J. Parrow, Broadcast psi-calculi with an application to wireless protocols, Softw. Syst. Model. 14 (1) (2015) 201–216, https://doi.org/10.1007/s10270-013-0375-z.
[41] J. Godskesen, S. Nanz, Mobility models and behavioural equivalence for wireless networks, in: COORDINATION, in: Lect. Notes Comput. Sci., vol. 5521, Springer, 2009, pp. 106–122.
[42] R. Vigo, F. Nielson, H. Nielson Broadcast, Denial-of-service, and secure communication, in: IFM, in: Lect. Notes Comput. Sci., vol. 7940, Springer, 2013, pp. 412–427.
[43] X. Wu, H. Zhu, Formal analysis of a calculus for WSNs from quality perspective, Sci. Comput. Program. (2018), https://doi.org/10.1016/j.scico.2017.08.007, in press.
[44] P. Attar, I. Castellani, Fine-grained and coarse-grained reactive noninterference, in: TGC, in: Lect. Notes Comput. Sci., vol. 8358, Springer, 2013, pp. 159–179.
[45] R. De Nicola, M. Loreti, R. Pugliese, F. Tiezzi, A formal approach to autonomic systems programming: the SCEL language, ACM Trans. Auton. Adapt. Syst. 9 (2014) 1–29, https://doi.org/10.1145/2619998.

Please cite this article in press as: R. Lanotte, M. Merro, A semantic theory of the Internet of Things, Inf. Comput. (2018), https://doi.org/10.1016/j.ic.2018.01.001