



ALMA MATER STUDIORUM
UNIVERSITÀ DI BOLOGNA

ARCHIVIO ISTITUZIONALE DELLA RICERCA

Alma Mater Studiorum Università di Bologna Archivio istituzionale della ricerca

Distributed runtime verification by past-CTL and the field calculus

This is the final peer-reviewed author's accepted manuscript (postprint) of the following publication:

Published Version:

Giorgio Audrito, Ferruccio Damiani, Volker Stolz, Gianluca Torta, Mirko Viroli (2022). Distributed runtime verification by past-CTL and the field calculus. THE JOURNAL OF SYSTEMS AND SOFTWARE, 187, 1-15 [10.1016/j.jss.2022.111251].

Availability:

This version is available at: <https://hdl.handle.net/11585/902468> since: 2022-11-14

Published:

DOI: <http://doi.org/10.1016/j.jss.2022.111251>

Terms of use:

Some rights reserved. The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. For all terms of use and more information see the publisher's website.

This item was downloaded from IRIS Università di Bologna (<https://cris.unibo.it/>).
When citing, please refer to the published version.

(Article begins on next page)

This is the final peer-reviewed accepted manuscript of:

Giorgio Audrito, Ferruccio Damiani, Volker Stolz, Gianluca Torta, Mirko Viroli, Distributed runtime verification by past-CTL and the field calculus, Journal of Systems and Software, Volume 187, 2022, 111251, ISSN 0164-1212.

The final published version is available online at: <https://doi.org/10.1016/j.jss.2022.111251>

Terms of use:

Some rights reserved. The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. For all terms of use and more information see the publisher's website.

This item was downloaded from IRIS Università di Bologna (<https://cris.unibo.it/>)

When citing, please refer to the published version.

Distributed Runtime Verification by Past-CTL and the Field Calculus

Giorgio Audrito^a, Ferruccio Damiani^a, Volker Stolz^b, Gianluca Torta^a, Mirko Viroli^c

^a *Università di Torino, Turin, Italy*

{giorgio.audrito,ferruccio.damiani,gianluca.torta}@unito.it

^b *Western Norway University of Applied Sciences, Bergen, Norway*
volker.stolz@hvl.no

^c *Alma Mater Studiorum—Università di Bologna, Cesena, Italy*
mirko.viroli@unibo.it

Abstract

Recent trends in the engineering of software-intensive systems increasingly promote the adoption of computation at the edge of the network, in the proximity of where sensing and actuation are performed. Applications are executed directly in IoT devices deployed in the physical environment, possibly with the aid of edge servers: there, interactions are essentially based on physical proximity, and communication with the cloud is sporadic if not absent.

The challenge of monitoring the execution of such system, by relying on local interactions only, naturally arises. We address this challenge by proposing a rigorous approach to distributed runtime monitoring for space-based networks of devices. We introduce the past-CTL logic, an extension of past-LTL able to express a variety of properties concerning the knowable past of an event. We formally define a procedure to derive, from a past-CTL formula, monitors that can be distributed on each device and whose collective behaviour verifies the validity of the formula at runtime across space and time. This is achieved by relying on the field calculus, a core programming language used to specify the behaviour of a collection of devices by viewing them as an aggregate computing machine, carrying out altogether a distributed computational process. The field calculus is shown to be a convenient language for our goals, since its functional composition approach provides a natural way of translating in a syntax-directed way properties expressed in a given logic into monitors for such properties. We show that the monitor process executing in each single device runs using local memory, message size, and computation time that are all linear in the size of the formula (1 bit per temporal connective). This matches the efficiency of the best available previous results for (non-distributed) monitors derived from past-LTL formulas. Finally, we empirically evaluate the applicability of the approach to sample problems in distributed computing, through simulated experiments with monitors written through a C++ library implementing the field calculus programming constructs.

Keywords: distributed systems, runtime verification, field calculus, temporal logic

2010 MSC: 00-01, 99-00

1. Introduction

Runtime verification is a computing analysis paradigm based on observing a system at runtime (to check its expected behaviour) by means of monitors generated from formal specifications [1]. Distributed runtime verification is runtime verification in connection with distributed systems: it comprises both monitoring of distributed systems and usage of distributed systems for monitoring. Being a verification technique, additionally, runtime verification promotes the generation of monitors from formal specifications, so as to precisely state the properties to check as well as providing formal guarantees about the results of monitoring. Distribution is hence a particularly challenging context in verification, for it requires to correctly deal with aspects such as synchronisation, faults in communications, possible lack of unique global time, and so on [2]. Additionally, the distributed system whose behaviour is to be verified at runtime could emerge from modern application scenarios like the Internet-of-Things (IoT), Cyber-Physical Systems (CPS), Edge Computing (EC) or large-scale Wireless Sensor Networks (WSN). In this case additional features are to be considered, like openness (the set of nodes is dynamic), large-scale (a monitoring strategy may need to scale from few units up to thousands of devices), and interaction locality (nodes may be able to communicate only with a small neighbourhood). So, in the most general case, distributed runtime verification challenges the way in which one can express properties on such dynamic distributed systems, can express flexible computational tasks, and can reason about compliance of properties and corresponding monitoring behaviour.

In this paper we considered distributed runtime verification for systems where a significant portion of computations is performed at the edge of the network, namely, where computational devices have a (possibly changing) physical position and space, and interact locally, with devices in the proximity. As a notable example, used throughout the paper (among others thoroughly discussed in Section 4), we will refer to a synthetic logical network comprising nodes of different nature (cloud, fog, and edge), sporadically issuing requests (of a few types) and then waiting for a corresponding response. In this context, we will show how significant properties of the system can be expressed and monitored.

Developing on our preliminary findings reported in [3], in this paper we argue that a promising approach to address the challenges of distributed runtime verification for these kinds of system [4] can be rooted on the computational paradigm of aggregate computing [5], along with the field calculus language [6]. Aggregate computing promotes a view of a distributed system as a conceptually single computing device, spread throughout the (physical or virtual) space in which nodes are deployed, that has been shown to be a good match for the IoT.

At the paradigm level, aggregate programming promotes the specification (construction, reasoning, programming) of global-level computational behaviour, where message exchange across individuals is essentially abstracted away through a declarative style to express interactions. At the modelling level, the field calculus can be leveraged, which expresses computations as transformations of *computational fields* (or *fields*

35 for short), namely, space-time distributed data structures mapping computational events (occurring at a given position of space and time) to computational values (the results of computation at that event). As an example, a set of temperature sensors spread over a building forms a field of temperature values (a field of reals), and a monitor alerting areas where the temperature was above a threshold for the last 10 minutes is a function from the temperature field to a field of booleans.

40 We see aggregate computing also as suitable for heterogeneous architectures such as edge computing, where both physical and logical neighbourhood relationships may hold between the terminal (IoT) devices and edge devices, and within or between the other network tiers beyond the edge up to a potentially centralised cloud instance. Its features naturally address challenges that are particular to such contexts [7]: mobility of devices, decentralization and fault tolerance, and interoperability and heterogeneity.

45 First, aggregate computing copes by design with mobility: On the one hand, it is agnostic to changing network performance and changing neighbours; On the other hand, it can easily be made aware of such notion by accessing sensors on performance data or including timestamps and clock-synchronization explicitly in all or some messages. Second, aggregate computing is also by nature decentralised: there is no central authority or aggregator per se, and any aggregation requires explicit design on top of broadcast-based communication: the absence of point-to-point connections and an abstraction from actual message transmission leads to a fault tolerant design by construction, as missing or delayed messages or network partitioning are indistinguishable from normal operation. Finally, interoperability and heterogeneity across platforms is — as in many software systems — solved by abstracting the communication layer.

55 The convenience of aggregate computing for distributed monitoring lies then in its compositionality mechanism, that is fully functional. Hence it induces a compositional translation (the translation of a phrase is determined by combining the translations of its subphrases) of properties expressed in a given logic into a fully-distributed field calculus program whose execution is actually a distributed monitor for such property. Most specifically, phrases are replicated across all nodes, hence an instance of the property is checked by every node of the system altogether.

60 In this paper we focus on past-CTL logic, a logic able to express properties concerning the knowable past of an event. Such properties include the eventual behaviour of participants in the network, interpreted based on which information has been gathered into an event from the past, and this is shown, through a variety of examples, to give forms of predictive capability. We show that the derived monitors match the optimal efficiency of previous results for past-LTL in a non-distributed setting [8]. Hence, this translation showcases how the field calculus can support distributed verification in a high-level compositional way, and paves the way for integration of future extensions of the logic.

The proposed approach aims to target highly-distributed systems where cloud-like support is sporadic or not available, or cannot be successfully exploited, making the monitoring of formulas with future modalities infeasible. In such contexts, using field calculus programming (which comes with fully implemented

70 tools such as FCPP [9] <http://fcpp.github.io>, a C++ internal DSL and simulator, PROTELIS [10] <http://protelis.github.io>, a JAVA external DSL, SCAFI [11] <https://scafi.github.io>, a SCALA internal DSL), in combination with past-CTL-based distributed runtime verification as proposed in this paper, can still allow to reasonably monitor interesting properties of distributed systems, as we show by examples and simulations. The main contribution of the present paper is twofold:

- 75 • The introduction of past-CTL, a modal logic for expressing temporal properties of a distributed system, inspired by the well-known temporal logics CTL and past-LTL.
- The proposal of a formal, detailed method to automatically translate formulas of past-CTL into fully distributed monitoring systems expressed by aggregate computing, which can be practically deployed and executed.

80 The remainder of this paper is organised as follows: Sect. 2 provides the necessary background on field calculus; Sect. 3 illustrates how the field calculus can be used to implement distributed monitors; Sect. 4 applies our novel solution to sample problems in aggregate computing; Sect. 5 discusses related work and Sect. 6 concludes.

2. A Recollection of the Field Calculus

85 The *field calculus* [6] is a minimal universal [12] language to express aggregate computations over distributed networks of (mobile) devices, each asynchronously capable of performing simple local computations and interacting with a neighbourhood by local exchanges of messages. Field calculus provides the necessary mechanism to express and functionally compose such distributed computations, by a level of abstraction that intentionally neglects explicit management of synchronisation, message exchanges between devices, position and quantity of devices, and so on—these aspects are however dealt with “under the hood”, namely, at the operational semantics level. In the field calculus, a program P is periodically and asynchronously executed on every device, according to a cyclic schedule executed by each device δ with period T_δ :

1. the device perceives contextual information formed by data provided by sensors, local information stored in the previous round, and messages collected from neighbours since the previous round (older messages get discarded after a certain timeout), the latter in the form of a *neighbouring value* ϕ —essentially a map from neighbour device identifiers δ to values ℓ ;
 2. when a round starts, the device evaluates the program P considering as input the contextual information gathered as described above (as formalised by the operational semantics);
 3. the result of local computation is a data structure that is stored locally, is broadcast to neighbours, and produces output values (possibly fed to actuators).
- 100

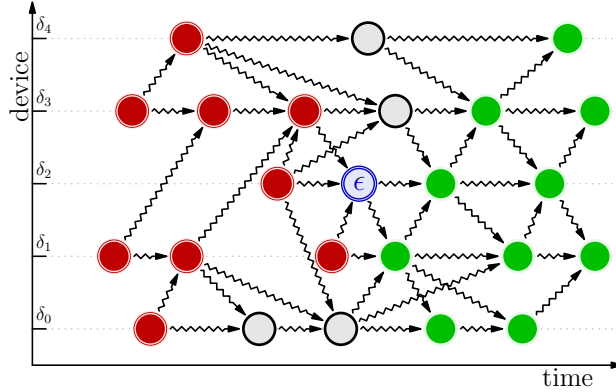


Figure 1: Sample event structure split into events in the causal past of ϵ ($\epsilon' < \epsilon$, circled solid red), events in the causal future ($\epsilon < \epsilon'$, solid green) and concurrent (non-ordered, hollow black). Events are sorted by the device on which they occur.

By repetitive execution of such computation rounds, across space (where devices are located) and time (when devices execute), a global behaviour emerges [13] at the overall network-level of interconnected devices, modelled as a single aggregate machine equipped with a neighbouring-based topology relation. Typically, the neighbouring relation reflects spatial proximity, but it could also be a logical relationship, e.g., connecting master devices to slave devices independently of their position, and the like. The data abstraction manipulated by the “aggregate computing machine” is hence a whole distributed *space-time value* Φ , which maps individual computation *events* ϵ (space-time points where and when a device sends messages at the end of a computation round) to data values v . The set of events, constituting the domain of fields, is then structured according to a *neighbouring* notion, dictating when an event can influence (by message-passing) another.

2.1. Event Structures

The causal relationship between events may then be formalised by the classical notion of *event structure* [14], which we will use in the following sections to interpret temporal logic formulas about distributed computations.

Definition 2.1 (Event Structure). An *event structure* $\mathbf{E} = \langle E, \rightsquigarrow, < \rangle$ is a countable set of events E together with a *neighbouring relation* $\rightsquigarrow \subseteq E \times E$ and a *causality relation* $< \subseteq E \times E$, such that the transitive closure of \rightsquigarrow forms the irreflexive partial order $<$ and the set $\{\epsilon' \in E \mid \epsilon' < \epsilon\}$ is finite for each ϵ (i.e., $<$ is *locally finite*).¹

Note that the transitive closure condition on \rightsquigarrow also implies that \rightsquigarrow is asymmetric and irreflexive. Figure 1 provides a sample event structure showing how, given an event ϵ , these relations partition events

¹The definition of event structure is usually given just in terms of the causality relation. We have also included the neighbouring relation since it is able to capture message passing details, which are needed to interpret most distributed programs.

into “causal past” (events from which information can potentially be carried to ϵ in a message), “causal future” (those to which information from ϵ can be carried) and “concurrent” (events informationally isolated from ϵ) subsets.

Since $<$ is uniquely induced by \rightsquigarrow , we omit it whenever convenient, or use its weak form \leq .² Notice that since $<$ is required to be irreflexive, \rightsquigarrow has to be an acyclic relation, thus inducing a directed acyclic graph (DAG) structure on E . In fact, \mathbf{E} can be thought of as a DAG with a “neighbouring” relation (modelling message passing) and a “reachability” relation (modelling causal dependence).

Event structures completely abstract away from the information about which device or devices might be performing an actual computation at each event, focussing on which data may be available at every computational step, no matter on what device the computation may be happening. Thus a series of computations on the same device (whether it is fixed or mobile) are to be modelled by a sequence of events $\epsilon_1, \dots, \epsilon_n$ such that $\epsilon_i \rightsquigarrow \epsilon_{i+1}$, in which message passing is implemented simply by keeping data available on the device for subsequent computations.

We do not make any assumptions on how a particular event structure come into being: each device may for example have changed its spatial position through actuators or external factors (which may have been reflected in particular position- or accelerator-sensor readings on its timeline), derived values may have evolved over time due to computation and communication, which could, but does not have to, have been implemented through field calculus programs.

Any sequence of computation events and message exchanges between them can be represented as an event structure, however, not all event structures are physically realisable by a distributed system following the computational model described at the beginning of this section. The subset of realisable event structures is characterised by the following definition.

Definition 2.2 (LUIC Augmented Event Structure). An *augmented event structure* is a tuple $\mathbb{E} = \langle E, \rightsquigarrow, <, d \rangle$ such that $\langle E, \rightsquigarrow, < \rangle$ is an event structure and $d : E \rightarrow D$ is a mapping from events to the devices where they happened. We define:

- $\text{next} : E \rightarrow E$ as the partial function³ mapping an event ϵ to the unique event $\text{next}(\epsilon)$ such that $\epsilon \rightsquigarrow \text{next}(\epsilon)$ and $d(\epsilon) = d(\text{next}(\epsilon))$, if such an event exists and is unique (i.e., $\text{next}(\epsilon)$ is the computation performed immediately after ϵ on the same device $d(\epsilon)$); and
- $\text{--}\rightarrow \subseteq E \times E$ as the relation such that $\epsilon \text{--}\rightarrow \epsilon'$ (ϵ *implicitly precedes* ϵ') if and only if $\epsilon' \rightsquigarrow \text{next}(\epsilon)$ and $\epsilon' \not\rightsquigarrow \epsilon$.

We say that \mathbb{E} is a *LUIC augmented event structure* if the following coherence constraints are satisfied:

²The weak form of a partial order is defined as $x \leq y$ iff $x < y$ or $x = y$.

³With $A \rightarrow B$ we denote the space of partial functions from A into B .

- **Linearity:** if $\epsilon \rightsquigarrow \epsilon_i$ for $i = 1, 2$ and $d(\epsilon) = d(\epsilon_1) = d(\epsilon_2)$, then $\epsilon_1 = \epsilon_2 = \text{next}(\epsilon)$ (i.e., every event ϵ is a neighbour of at most another one on the same device);
- **Uniqueness:** if $\epsilon_i \rightsquigarrow \epsilon$ for $i = 1, 2$ and $d(\epsilon_1) = d(\epsilon_2)$, then $\epsilon_1 = \epsilon_2$ (i.e., neighbours of an event all happened on different devices);
- **Impersistence:** if $\epsilon \rightsquigarrow \epsilon_i$ for $i = 1, 2$ and $d(\epsilon_1) = d(\epsilon_2) = \delta$, then either $\epsilon_2 = \text{next}^n(\epsilon_1)$ and $\epsilon \rightsquigarrow \text{next}^k(\epsilon_1)$ for all $k \leq n$, or the same happens swapping ϵ_1 with ϵ_2 (i.e., an event reaches a contiguous set of events on a same device);
- **Computation immediacy:** the relation $\rightsquigarrow \cup \dashrightarrow$ is acyclic on E (i.e., explicit causal dependencies $<$ are consistent with implicit time dependencies \dashrightarrow).

The first two constraints are necessary for defining the semantics of an aggregate program (denotational semantics in [6, 15]). The third reflects that messages are not retrieved after they are first dropped (and in particular, they are all dropped on device reboots). The last constraint reflects the assumption that computation and communication are modelled as happening instantaneously. In this scenario, the explicit causal dependencies imply additional time dependencies $\epsilon \dashrightarrow \epsilon'$: if ϵ' was able to reach $\text{next}(\epsilon)$ but not ϵ , the firing of ϵ' must have happened *after* ϵ .

The event structure in Figure 1 satisfies the LUIC constraints with the represented device assignment. Interpreting this structure in terms of physical devices and message passing, a physical device is instantiated as a chain of events connected by \rightsquigarrow relations (representing evolution of state over time with the device carrying state from one event to the next), and any \rightsquigarrow relation between devices represents information exchange from the tail neighbour to the head neighbour. As a matter of fact, this is a very flexible and permissive model: there are no assumptions about synchronisation, shared identifiers or clocks, or even regularity of events (though of course these things are not prohibited either).

Though rather abstract, the notion of (augmented) event structure is well-suited to ground a semantics for *space-time computations*, intended as “elaborations of distributed data in a network of related events”: the causality ordering of events abstracts time, while the presence of concurrent events abstract spatial dislocation. We refer to [6] for the definition of a denotational semantics of the field calculus based on augmented event structures and for a formal account of its relation with the operational semantics of the field calculus presented below.

Definition 2.3 (Space-Time Values). Let \mathbf{V} be a denumerable universe of allowed computational values and \mathbb{E} be a given LUIC augmented event structure. A *space-time value* Φ in \mathbb{E} is an annotation of the graph \mathbb{E} with labels in \mathbf{V} , that is, a tuple $\Phi = \langle \mathbb{E}, f \rangle$ with $f : E \rightarrow \mathbf{V}$, taking E as the set of events in \mathbb{E} .

Space-time values model data *spatially* distributed across devices and *temporally* distributed across time: in this way, time-evolving inputs, sensor information (e.g., time clock, temperature) and intermediate results

185 of computations (which are naturally time-dependent) are easily represented, attaining maximal generality while ensuring composability of behaviour. These quantities can be manipulated by distributed computations (i.e., consumed as inputs) and can also be created by them (i.e., produced as outputs). Thus, an aggregate computer is a “collective” device manipulating such space-time values, modelled as a *space-time function*.

Definition 2.4 (Space-Time Function). Let $\mathbf{V}(\mathbb{E}) = \{\langle \mathbb{E}, f \rangle \mid f : E \rightarrow \mathbf{V}\}$ be the set of all possible space-time values in an augmented event structure \mathbb{E} . Then, an *n-ary space-time function in \mathbb{E}* is a partial map
190 $\mathbf{f} : \mathbf{V}(\mathbb{E})^n \rightarrow \mathbf{V}(\mathbb{E})$.

The definition of a space-time function \mathbf{f} requires every input and output space-time value to exist in the same augmented event structure \mathbb{E} . However, it does not specify how the output space-time values are obtained from the inputs, and in fact not all space-time functions \mathbf{f} are physically realisable by a program,
195 as \mathbf{f} may violate either causality or Turing-computability (see [12] for further details).

The specification of a space-time function can be either done at a low-level (i.e. through local interactions), in order to define programming language constructs and general-purpose building blocks of reusable behaviour, or at a high-level (i.e. by composition of other space-time functions with a global interpretation) in order to design collective adaptive services and whole distributed applications—which ultimately work by
200 getting input fields from sensors and process them to produce output fields to actuators. However, in aggregate computing a distributed program P always has both the *local* and *global* interpretations, dually linked: the global interpretation as a space-time function (obtained through a denotational semantics [6, 15]), and the local interpretation as a procedure performed in a firing (defined by the operational semantics presented below).

205 2.2. Syntax and Operational Semantics

The syntax of the fragment of the field calculus that fits the needs of this paper is given in Figure 2—we refer to [6, 16] for a complete presentation of the syntax, type system and operational semantics of field calculus: in the following, we present only what is strictly needed for the remainder of this paper. The overbar notation \bar{e} is a shorthand for sequences of elements, and multiple overbars are intended to be
210 expanded together, e.g., \bar{e} stands for the sequence of expressions e_1, \dots, e_n and $\bar{\delta} \mapsto \bar{\ell}$ stand for the map $\delta_1 \mapsto \ell_1, \dots, \delta_n \mapsto \ell_n$ (that associates device identifiers δ to local values ℓ). The keyword **share** is the main peculiar construct of the field calculus, responsible of both interaction and field dynamics; while **def** and **if** correspond to the standard function definition and the branching expression constructs.

A program P consists of a list of function definitions \bar{F} , each written as “**def** $d(x_1, \dots, x_n) \{e\}$ ”, followed
215 by a main expression e that is the one executed at each computation round (as well as the one representing the overall field computation, in the global viewpoint). An expression e can be:

- A *variable* x , used e.g. as formal parameter of functions.

$P ::= \bar{F} e$	program
$F ::= \text{def } d(\bar{x}) \{e\}$	function declaration
$e ::= x \mid f(\bar{e}) \mid v \mid \text{if}(e)\{e\}\{e\} \mid \text{share}(e)\{(x)=>e\}$	expression
$f ::= d \mid b$	function name
$v ::= \ell \mid \phi$	value
$\ell ::= c(\bar{\ell})$	local value
$\phi ::= \bar{\delta} \mapsto \bar{\ell}$	neighbouring field value

Figure 2: Syntax of the field calculus language.

- A *value* v , which can be of the following two kinds:
 - A *local value* ℓ , with structure $c(\bar{\ell})$ or simply c when $\bar{\ell}$ is empty (defined via data constructor c and arguments $\bar{\ell}$), can be, e.g., a Boolean (**true** or **false**), a number, a string, or a structured value (e.g., a pair **pair(true,5)**).
 - A *neighbouring (field) value* ϕ that associates neighbour device identifiers δ to local values ℓ , e.g., it could be the neighbouring value of distances to neighbours—note that neighbouring field values are not part of the surface syntax, they are produced at runtime by evaluating expressions, as described below.
- A function call $f(\bar{e})$, where f can be of two kinds: a *user-declared function* d (declared by the keyword **def**, as illustrated above) or a *built-in function* b , such as a mathematical or logical operator, a data structure operation, or a function returning the value of a sensor.
- A branching expression $\text{if}(e_1)\{e_2\}\{e_3\}$, used to split a field computation in two isolated sub-networks, where/when e_1 evaluates to **true** or **false**: the result is computation of e_2 in the former area, and e_3 in the latter.
- A share expression $e = \text{share}(e_1)\{(x)=>e_2\}$, which incorporates message passing and local state evolution. The result of such expression is obtained by:
 - Gathering the results $\bar{\ell}$ obtained by neighbours $\bar{\delta}$ for the whole expression e in their last rounds into a neighbouring field value $\phi = \bar{\delta} \mapsto \bar{\ell}$.
 - Such ϕ may also contain a value for the current device $\delta = \delta_i$ for some i (if it is not the first round it executes e). If not, the result ℓ of evaluating e_1 is used as value for the current device δ and incorporated into $\phi := \bar{\delta} \mapsto \bar{\ell}, \delta \mapsto \ell$.

- Expression e_2 is then evaluated by substituting ϕ to x , obtaining the overall value ℓ' for e .
- Value ℓ' is broadcast to neighbours, allowing them to use it in constructing their following neighbours' observation ϕ as described above.

This construct is designed for structuring device interaction, but can also be used for evolving a state locally, if the variable x appears in e_2 within the built-in operator `locHood(ϕ)`, which extracts the value $\phi(\delta)$ relative to the current device δ from a neighbouring field value ϕ . Note that the evaluation by a device δ of a `share`-expression within a branch of some `if(e_1)` expression, is affected only by the neighbours of δ that, during their last computation cycle, evaluated the same value for the guard e_1 , since these are the only ones that computed that share expression thus sharing a value for it with δ .

Example 2.5. As an example illustrating the constructs of field calculus, first consider the expression:

```
share (false) { (old) => f || anyHood(old) }
```

where f is some captured variable, and `anyHood` is a built-in operator collapsing a field $\phi = \bar{\delta} \mapsto \bar{b}$ of Boolean values \bar{b} into the disjunction of its constituent values $\bigvee \bar{b}$. When this expression is evaluated for the first time, no messages are available from neighbours, hence the gathered ϕ is comprised only of the value of e_1 for the current device: $\phi = \delta \mapsto \text{false}$. As ϕ is substituted for `old`, we compute `anyHood(ϕ)` which is false, and finally `f || false` which is f . When this same expression is evaluated later on, ϕ may contain messages from neighbours. If any one of them is `true`, `anyHood(ϕ)` will be true and so will the whole share expression. If none of them is true, the value of the whole expression will be again equal to f . This distributed behaviour can be understood as computing whether f has *ever* been true by *gossiping* whether this is the case across neighbours. As such, this expression can be conveniently encapsulated into the following function definition:

```
def gossip-ever(f) { share (false) { (old) => f || anyHood(old) } }
```

As discussed above, function `gossip-ever` takes a boolean field f and whenever it holds a `true` value in a device, this gets propagated throughout the network by gossiping.

The built-in functions and data constructors used in this paper are listed in Figure 3 together with their types and formal interpretations—all of these operators are natively available in existing implementations of the field calculus, including FCPP (used in the case studies presented in Section 4). In this paper, we use types `bool` (Boolean values), `field[bool]` for neighbouring fields built from values of type `bool`, `pair[bool, bool]` for pairs of Boolean values, and types $(\bar{T}) \rightarrow T$ for functions. In the remainder of this paper, we also use $[e_1, e_2]$ as short for `pair(e_1, e_2)`.

Constructors:		
<code>true, false</code>	$= () \rightarrow \text{bool}$	\top, \perp
Built-ins:		
<code>!</code>	$= (\text{bool}) \rightarrow \text{bool}$	\neg
<code> , &&</code>	$= (\text{bool}, \text{bool}) \rightarrow \text{bool}$	\vee, \wedge
<code><=, ==</code>	$= (\text{bool}, \text{bool}) \rightarrow \text{bool}$	$\leq, =$
<code>pair</code>	$= (\text{bool}, \text{bool}) \rightarrow \text{pair}[\text{bool}, \text{bool}]$	$(p, q) \mapsto [p, q]$
<code>fst, snd</code>	$= \text{pair}[\text{bool}, \text{bool}] \rightarrow \text{bool}$	$[p, q] \mapsto p$ (resp. q)
<code>locHood</code>	$= (\text{field}[\text{bool}]) \rightarrow \text{bool}$	$\phi \mapsto \phi(\delta)$
<code>anyHood</code>	$= (\text{field}[\text{bool}]) \rightarrow \text{bool}$	$\phi \mapsto \bigvee \{\phi(\delta') \mid \delta' \in \mathbf{dom}(\phi)\}$
<code>allHood</code>	$= (\text{field}[\text{bool}]) \rightarrow \text{bool}$	$\phi \mapsto \bigwedge \{\phi(\delta') \mid \delta' \in \mathbf{dom}(\phi)\}$

Figure 3: Types and interpretations of the data constructors and built-in functions used throughout this paper. We use δ to denote the current device of the current event, and notation $\mathbf{dom}(\phi)$ to denote the domain $\bar{\delta}$ of a function $\phi = \bar{\delta} \mapsto \bar{\ell}$.

$\psi ::= \perp \mid \top \mid q \mid (\neg\psi) \mid (\psi \wedge \psi) \mid (\psi \vee \psi) \mid (\psi \Rightarrow \psi) \mid (\psi \Leftrightarrow \psi)$	logical op.
$\mid (Y\psi) \mid (AY\psi) \mid (EY\psi) \mid (\psi S\psi) \mid (\psi AS\psi) \mid (\psi ES\psi)$	temporal op.
$\mid (P\psi) \mid (AP\psi) \mid (EP\psi) \mid (H\psi) \mid (AH\psi) \mid (EH\psi)$	

Figure 4: Syntax of past-CTL.

3. Monitoring of Past-CTL Properties in Field Calculus

In this section, we introduce a temporal logic (past-CTL) designed for reasoning about properties of (augmented) event structures, which can be translated in field calculus in order to allow efficient recursive computation of the truth value of such formulas, assuming that every participant is evaluating the same property from their perspective with regard to any quantifiers. This logic will be based on atomic propositions (observables in a particular state), which we assume to be backed by implicit Boolean space-time values. This will allow us not to consider a mix of temporal operators and (boolean) field calculus expressions, but rather completely encapsulate the latter in the abstract propositions.

Sect. 3.1 introduces the syntax of past-CTL, while Sect. 3.2 provides its formal interpretation on an augmented event structure. Sect. 3.3 introduces a translation of past-CTL into field calculus, coherently with the formal interpretation, and Sect. 3.4 concludes discussing the applicability of past-CTL and its limited future prediction capabilities.

3.1. Past-CTL Syntax

Figure 4 presents the past-CTL syntax, following the syntax of past-LTL operators in literature (see e.g. [17]) and extending them as CTL extends LTL. This logic is almost identical to traditional CTL, with

two main differences:

- temporal operators are interpreted in the past (and thus their names are changed accordingly), along paths of message exchanges in an event structure, composed of events that all happened in the past (and are not alternative realities);
- there are un-quantified versions of the operators along with quantified versions, which refer to the linear past on a same device (and thus behave as past-LTL operators).

Formulas in past-CTL can be interpreted in augmented event structures (see Sect. 2.1), giving a definite truth value for each event (and thus producing a space-time value), as we shall see in Sect. 3.2. The operators have the following informal meaning:

- similarly to CTL, path quantifiers A and E refer to the fact that the corresponding formula holds in every (resp. some) path of messages, starting from an initial event of a device and ending in the current event; whereas the absence of path quantifiers refers to the fact that the corresponding formula holds in the special path of messages connecting the events happened in the current device;
- operator Y “*yesterday*” postulates that ψ held in the previous event on the considered path (on the same device for Y, on every other device for AY, in another device for EY);
- operator S “*since*” postulates that its second argument held in some past event in the considered path, and its first argument has held since then;
- operator P “*previously*” postulates that ψ held in some past event on the considered path;
- operator H “*historically*” postulates that ψ held in every past event on the considered path.

The list of operators presented above is redundant. A minimal set of operators is $\neg, \vee, Y, AY, S, AS, ES$, the others being expressible as:

- $EY \psi := \neg AY \neg \psi$;
- $P \psi := \top S \psi$ (similarly for AP, EP with AS, ES);
- $H \psi := \neg P \neg \psi$ (similarly for AH, EH with EP, AP).

Example 3.1. $(AH(r \Rightarrow (Y(\neg r S q))))$ is a valid past-CTL formula, with the informal meaning “it is always and everywhere the case, that if r holds then in the previous round r was false since a past round where q was true”. In the following we shall apply usual operator precedence conventions (assuming temporal operators to bind tighter than logical ones), in order to write the same formula as $AH(r \Rightarrow Y(\neg r S q))$.

Referring to the scenario on request/response matching over a logical layered network (described in the introduction), we can imagine r to correspond to a *response* and q to correspond to a *request*. In this

context, the formula can be interpreted as “every response corresponds to a preceding request”: in particular, $r \Rightarrow Y(\neg r S q)$ holds in an event with a recent trace of the kind $r, \neg r, \dots, \neg r, q$.

3.2. Semantic Interpretation on Augmented Event Structures

320 Given a past-CTL formula ψ , an augmented event structure \mathbb{E} , and a Boolean space-time value Φ_q in it for each propositional symbol q occurring in ψ , we can interpret ψ in \mathbb{E} obtaining a space-time value $\Phi = \mathbb{E}[\psi]$, such that $\mathbb{E}[\psi](\epsilon)$ is defined by causal recursion on ϵ as follows.

- $\mathbb{E}[\top](\epsilon) = \top$, and $\mathbb{E}[q](\epsilon)$ is $\Phi_q(\epsilon)$;
- $\mathbb{E}[\neg\psi](\epsilon) = \neg\mathbb{E}[\psi](\epsilon)$ and $\mathbb{E}[\psi_1 \vee \psi_2](\epsilon) = \mathbb{E}[\psi_1](\epsilon) \vee \mathbb{E}[\psi_2](\epsilon)$;
- 325 • $\mathbb{E}[Y \psi](\epsilon) = \mathbb{E}[\psi](\epsilon')$ where ϵ' is the event preceding ϵ on the same device (if it exists, $\mathbb{E}[Y \psi](\epsilon) = \perp$ otherwise);
- $\mathbb{E}[AY \psi](\epsilon) = \bigwedge_{\epsilon' \rightsquigarrow \epsilon} \mathbb{E}[\psi](\epsilon')$ (i.e., ψ is true in each preceding event ϵ');
- $\mathbb{E}[\psi_1 AS \psi_2](\epsilon)$ holds iff for every path $\epsilon_1 \rightsquigarrow \dots \rightsquigarrow \epsilon_n = \epsilon$ such that ϵ_1 has no neighbours, $\mathbb{E}[\psi_2](\epsilon_i)$ holds for some i and $\mathbb{E}[\psi_1](\epsilon_j)$ holds for each $j = i + 1 \dots n$;
- 330 • $\mathbb{E}[\psi_1 ES \psi_2](\epsilon)$ holds iff it exists a path $\epsilon_1 \rightsquigarrow \dots \rightsquigarrow \epsilon_n = \epsilon$ such that $\mathbb{E}[\psi_2](\epsilon_1)$ holds and $\mathbb{E}[\psi_1](\epsilon_i)$ holds for $i = 2 \dots n$;
- $\mathbb{E}[\psi_1 S \psi_2](\epsilon)$ holds iff it exists a path $\epsilon_1 \rightsquigarrow \dots \rightsquigarrow \epsilon_n = \epsilon$ of events all occurring on the same device $\delta = d(\epsilon)$,⁴ such that $\mathbb{E}[\psi_2](\epsilon_1)$ holds and $\mathbb{E}[\psi_1](\epsilon_i)$ holds for $i = 2 \dots n$.

Semantics of derived operators can be inferred accordingly.

Example 3.2. The semantics of formula $\psi := AH(r \Rightarrow Y(\neg r S q))$ is:

$$\begin{aligned}
\mathbb{E}[\psi](\epsilon) &= \bigwedge_{\epsilon' \leq \epsilon} \mathbb{E}[r \Rightarrow Y(\neg r S q)](\epsilon') = \bigwedge_{\epsilon' \leq \epsilon} [\Phi_r(\epsilon') \Rightarrow \mathbb{E}[Y(\neg r S q)](\epsilon')] \\
&= \bigwedge_{\epsilon' \leq \epsilon} \left[\Phi_r(\epsilon') \Rightarrow \bigvee_{\substack{\epsilon'' \rightsquigarrow \epsilon' \\ d(\epsilon'')=d(\epsilon')}} \mathbb{E}[\neg r S q](\epsilon'') \right] \\
&= \bigwedge_{\epsilon' \leq \epsilon} \left[\Phi_r(\epsilon') \Rightarrow \bigvee_{\substack{\epsilon'' \rightsquigarrow \epsilon' \\ d(\epsilon'')=d(\epsilon')}} \bigvee_{\substack{\epsilon_1 \rightsquigarrow \epsilon_2 \dots \rightsquigarrow \epsilon_n = \epsilon'' \\ d(\epsilon_i)=d(\epsilon_j) \forall i,j}} \left(\Phi_q(\epsilon_1) \wedge \bigwedge_{i \geq 2} \neg \Phi_r(\epsilon_i) \right) \right]
\end{aligned}$$

```

def Y(f) { snd(share([false,false]){(old) => [f, locHood(fst(old))])} }
def AY(f) { snd(share([true, true ]){(old) => [f, allHood(fst(old))])} }
def EY(f) { snd(share([false,false]){(old) => [f, anyHood(fst(old))])} }
def S(f1, f2) { share (false) {(old) => f2 || (f1 && locHood(old))} }
def AS(f1, f2) { share (false) {(old) => f2 || (f1 && allHood(old))} }
def ES(f1, f2) { share (false) {(old) => f2 || (f1 && anyHood(old))} }
def P(f) { share (false) {(old) => f || locHood(old)} }
def AP(f) { share (false) {(old) => f || allHood(old)} }
def EP(f) { share (false) {(old) => f || anyHood(old)} }
def H(f) { share (true) {(old) => f && locHood(old)} }
def AH(f) { share (true) {(old) => f && allHood(old)} }
def EH(f) { share (true) {(old) => f && anyHood(old)} }

```

Figure 5: Encoding of past-CTL operators as field calculus functions.

335 3.3. Automatic Translation in Field Calculus

Past-CTL operators can be encoded as field calculus functions, taking the Boolean value of their subformulas as arguments. Figure 5 shows a possible implementation for them, assuming that:

- `fst`, `snd` are operators extracting the first and second element of a pair;
- `anyHood` is as in Example 2.5;
- 340 • `allHood` is a built-in operator collapsing a field ϕ of Boolean values into the conjunction of its constituent values $\bigwedge \phi$;
- `locHood` is a built-in operator collapsing a field ϕ into the value $\phi(\delta)$ held by ϕ for the current device δ .

Using these functions, past-CTL formulas can be straightforwardly translated into field calculus, as shown 345 in Figure 6. We translate atomic propositions q into built-in function calls getting their value from some external environment. We translate logical operators into their corresponding built-ins, assuming that `false` < `true`, as in common programming languages such as C/C++ or Python. Temporal operators are translated by calling the corresponding functions.

Example 3.3. The translation of formula $AH(r \Rightarrow Y(\neg r S q))$ is $AH(r \leq Y(S(!r, q)))$.

⁴We recall that $d(\epsilon)$, according to the definition of augmented event structure, corresponds to the device where ϵ happened.

$\llbracket \top \rrbracket = \text{true}$	$\llbracket \psi_1 \vee \psi_2 \rrbracket = \llbracket \psi_1 \rrbracket \ \ \llbracket \psi_2 \rrbracket$
$\llbracket \perp \rrbracket = \text{false}$	$\llbracket \psi_1 \wedge \psi_2 \rrbracket = \llbracket \psi_1 \rrbracket \ \&\& \ \llbracket \psi_2 \rrbracket$
$\llbracket q \rrbracket = \mathbf{q}()$	$\llbracket \psi_1 \Rightarrow \psi_2 \rrbracket = \llbracket \psi_1 \rrbracket \ \lt;= \ \llbracket \psi_2 \rrbracket$
$\llbracket \neg \psi \rrbracket = \mathbf{!} \llbracket \psi \rrbracket$	$\llbracket \psi_1 \Leftrightarrow \psi_2 \rrbracket = \llbracket \psi_1 \rrbracket \ \mathbf{==} \ \llbracket \psi_2 \rrbracket$
$\llbracket Y \psi \rrbracket = \mathbf{Y}(\llbracket \psi \rrbracket)$	$\llbracket \psi_1 S \psi_2 \rrbracket = \mathbf{S}(\llbracket \psi_1 \rrbracket, \llbracket \psi_2 \rrbracket)$
$\llbracket AY \psi \rrbracket = \mathbf{AY}(\llbracket \psi \rrbracket)$	$\llbracket \psi_1 AS \psi_2 \rrbracket = \mathbf{AS}(\llbracket \psi_1 \rrbracket, \llbracket \psi_2 \rrbracket)$
$\llbracket EY \psi \rrbracket = \mathbf{EY}(\llbracket \psi \rrbracket)$	$\llbracket \psi_1 ES \psi_2 \rrbracket = \mathbf{ES}(\llbracket \psi_1 \rrbracket, \llbracket \psi_2 \rrbracket)$
$\llbracket P \psi \rrbracket = \mathbf{P}(\llbracket \psi \rrbracket)$	$\llbracket H \psi \rrbracket = \mathbf{H}(\llbracket \psi \rrbracket)$
$\llbracket AP \psi \rrbracket = \mathbf{AP}(\llbracket \psi \rrbracket)$	$\llbracket AH \psi \rrbracket = \mathbf{AH}(\llbracket \psi \rrbracket)$
$\llbracket EP \psi \rrbracket = \mathbf{EP}(\llbracket \psi \rrbracket)$	$\llbracket EH \psi \rrbracket = \mathbf{EH}(\llbracket \psi \rrbracket)$

Figure 6: Straightforward compositional translation of past-CTL into field calculus.

350 We recall that the field calculus programs obtained from past-CTL formulas can be straightforwardly encoded in one of the DSLs (FCPP, PROTELIS and SCAFI, mentioned in Sect. 1) that implement its programming model and then deployed and executed on a network of devices.

It is worth observing that the translation is fully compositional, and resembles the syntax of a non-distributed monitor, even though it is executed in a fully distributed manner thanks to the peculiar field calculus execution model and semantics. Compositionality also allows to easily handle extensions of the logic, as additional operators would just translate into additional rows in the translation table. Furthermore, computational complexity matches that of the best known monitors for past-LTL formulas in a non-distributed setting [8].

Theorem 3.4 (From past-CTL to field calculus). *The translation in Figure 6 is correct, computes in space/time which is linear in the size of the formula (in each event), and exchanges 1 bit for every S, P, H operator and 2 bits for every Y operator.*⁵

Proof. Let ψ be composed from sub-formulas ψ_1, \dots, ψ_n ($0 \leq n \leq 2$) and proceed by structural induction on formulas. The translation of ψ consists of a function call, where the body computes in constant space/time from the argument values, that are obtained from the translations of ψ_i which compute in linear space/time by inductive hypothesis. It follows that the overall space/time required is also linear. Furthermore, if ψ is not built from a temporal operator, its direct translation does not involve message-exchanging constructs. If it is built from a S, P, H operator, it consists of a single message-exchanging construct (here: **share**) sharing a single Boolean value (1 bit). If it is built from a Y operator, it also consists of a single message-exchanging construct, which however shares a pair of two Boolean values (2 bits).

⁵Optimizations in order to exchange 1 bit for every Y operator are possible, but require additional field calculus constructs, thus we decided to omit them for presentation simplicity.

370 If ψ is built from a traditional logic operator, correctness is trivial. If $\psi = Y \psi_1$, the translation of ψ calls a function, which takes the second element of a pair consisting of the current value $\mathbb{E}[\psi_1](\epsilon)$ of ψ_1 (by inductive hypothesis) together with the value the first element had in the previous event on the same device, that is the previous value of ψ_1 on the same device. If $\psi = AY \psi_1$, the translation of ψ is taking the second element of a pair which consists of $\mathbb{E}[\psi_1](\epsilon)$ together with the conjunction of the values $\mathbb{E}[\psi_1](\epsilon')$ the first
 375 element had in previous events. A similar argument applies for EY.

If $\psi = \psi_1 S \psi_2$ and $\mathbb{E}[\psi_2](\epsilon)$ is true, then $\mathbb{E}[\psi](\epsilon)$ is true and so is the translation. If it is not, then ψ_1 has to be true on previous events on the same device, since a preceding event where ψ_2 held. This is equivalent to asking that ψ_1 holds in the current event, and ψ held on the preceding event, which in field calculus corresponds to `f1 && locHood(old)`. Similar arguments apply for AS, ES. \square

380 *Remark 3.5.* Theorem 3.4 shows that the field calculus translation provided is asymptotically optimal, since some past-CTL formulas cannot be computed with sub-linear time or space, nor with message size lower than 1 bit per temporal connective. Such a family of formulas is, for example, $\bigwedge_{i \in N} EP q_i$ where q_i are independent atomic propositions. A computation in sub-linear time cannot possibly check every q_i and thus cannot be correct. Furthermore, it is necessary for every event to message whether each q_i has ever become
 385 true, otherwise some past information would be lost making it impossible to reconstruct the correct answer.

3.4. Practical Use of Persistent Past-CTL Properties

Past-CTL logic states properties about the past of a given event in space and time. Like traditional semantics of CTL, our semantics is defined for a state (event) by checking derived properties on preceding states and combining them into a verdict *for that state*. The decision to check a logical property for one
 390 state, some states or all states of an event structure is up to the end-user. However, as we are considering an event structure that is evolving at runtime and hence cannot make this choice explicitly, we re-interpret a formula as stating properties about *what we know so far of the future* of a system: our mechanism implicitly checks a given property on every state, conceptually “restarting” evaluation of the overall (initial) formula from each new state. We say that a formula is persistent in the sense that it is checked again and again,
 395 based on potentially new information.

A main concern for runtime verification are safety or liveness properties (“something bad never happens” / “something good eventually happens”), and through our choice of a past-time logics we have cemented our viewpoint into the past, clearly defining on which observations we want to base our evaluation. However, from the perspective of an agent starting up, taking such a property and (only) evaluating it in the initial
 400 state (like the root of a Kripke structure) where it does not have a past yet or did not yet receive any information from neighbours clearly does not make sense. Instead, we want to know the verdict for every state that an agent goes through.

Some of the properties will be monotonic, in the sense that an intermediate verdict fixes all future verdicts, either through quantification across neighbours (A/E) or over the past (P/H).

405 For instance, assume that we are interested in whether a certain property φ is ever true in any point of space and time. In an event infinitely far in the future, this could be checked through EP φ . The same formula can be interpreted in a “present” event as computing whether φ has already become true (\top result), or *not yet* (\perp result) based on the limited knowledge so far. Similarly, if we want to check whether a certain assumption φ is violated, we could monitor it through AH φ , interpreting its result on events as either *good*
 410 *so far* (\top) or *violation detected* (\perp).

The predictive capabilities of past-CTL formulas are mostly restricted to the interpretation of their outermost connective, and are thus more limited compared to those of a full-fledged future logic, which is sometimes able to perform inferences in order to detect events ahead of time. However, that increase in predictive capability comes with a higher computational cost; and in many practical scenarios the improve-
 415 ment is not large enough (if at all present) to justify the additional costs. In the following section, we will inspect some such examples.

3.5. Integration of Field Calculus Monitors and External Applications

Thus far, we have discussed how to express relevant properties of a distributed system in past-CTL and how to generate distributed Field Calculus programs for monitoring such properties. However, when the
 420 system to be monitored is itself made of software, or is wrapped in a software system, it becomes fundamental to integrate it with the generated monitors. In this section we briefly consider some lines for realizing such integration.

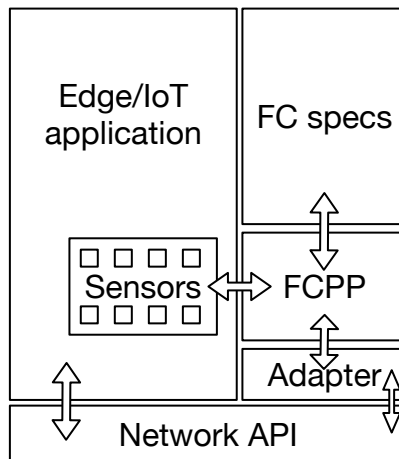


Figure 7: Architecture / Integration

Let us first consider a C/C++ distributed application. Fig. 7 gives a high-level view of how formulas

in FC are checked by the FCPP runtime in parallel to a running application. The FCPP runtime needs,
425 on one hand, to access to the sensor data wrapped in the application, and on the other hand an adapter
to access the network, either in the form of physical access to e.g. Ethernet, 5G, ..., or a software defined
network.

At the source code level, first of all the application must import the FCPP library and integrate the
generated formulas (as we have done for the sample applications described in Section 4). Notice that no
430 knowledge on aggregate computing is required: only an informal understanding of temporal logics, so that
the programmer can express the formulas he needs to monitor (everything else is fully automated). The
network adapter code is provided by implementing a transceiver interface for the underlying platform, i.e.
`os::uid` (the unique id of the devices on the network) and `os::transceiver` (functions to send and receive
data on the network). Then, the monitor code can be run as a *net* object in its separate thread as in the
435 following.

```
#include "lib/fcpp.hpp"

namespace fcpp {
440 namespace os {
    device_t uid() { ... } // unique identifier
    struct transceiver { ... }; // os interface for send-receive of packets
}
445
struct P {}; ... // propositional variables
struct M {}; ... // monitored formulas

MAIN() {
450     bool p = node.storage(P{}); ... // access propositions
     bool m = ... // compute monitors (past-CTL formulas)
     node.storage(M{}) = m; .. // write monitors' results
}

455 using net_type = component::monitor<
    component::tags::program<main>, component::tags::exports<bool>,
    component::tags::tuple_store<P, bool ... M, bool ...>
>::net;
460 }

int main() {
    using namespace fcpp;
    net_type network{common::make_tagged_tuple<>()};
465     std::thread monitor_thread(network.run);
    ... // main application code
}
```

When the main application code needs to provide sensor data to the monitor, it can access the `storage`

470 of the current monitor node to set the current values of the propositional variables that act as inputs in the property being monitored, and then release a new round of communication among the monitor nodes. Similarly, to access the current value of the monitored property, it is sufficient to read it from the node storage:

```
475 typename net_type::lock_type l;  
auto& node = net.node_at(os::uid(), l);  
node.storage(P{}) = ...  
node.trigger();  
m = node.storage(M{})  
480 ... // react to violations of the property
```

The integration of JVM-based applications (e.g., Java, Scala, Kotlin) with the past-CTL monitors is conceptually similar to the one just described for C++, except that the Scafi library should be used in place of FCPP.

485 Finally, a software system based on any technology can be integrated with FCPP monitors by writing an ad-hoc, small middleware C++ system that, on one hand, wraps the Field Calculus monitors and, on the other hand, exchanges messages with the monitored system exploiting inter-process communication mechanisms such as message queues, pipes, or even files.

4. Sample Applications

490 In this section, we present sample applications of past-CTL logic to the monitoring of distributed systems coming from various settings: service discovery (Sect. 4.1), swarm control (Sect. 4.2), disaster management (Sect. 4.3), and smart home scenarios (Sect. 4.4). All the simulations have been performed within the FCPP language and simulator [9] and are available online.⁶ In every scenario, we assumed the computation frequency of devices to be about $1Hz$ (with 10% standard deviation), and plotted the percentage of *false* values for monitored formulas across devices in the network over time. The scenarios involve both static sensors or logical devices, units on humans moving at a walking speed of $1.4m/s$ [18], and drones flying with a top speed of $15m/s$ ($54km/h$), which is an average limit for commercial UAVs.

4.1. Service Discovery

We consider a network with three types of nodes: *cloud* nodes located on a small central circle; *fog* nodes 500 located on an inner circle at some distance from the cloud; and *edge* nodes located on an outer circle further away from the cloud. Proposition rq_i is true iff the current node has just requested a service of type i , and proposition rs_i is true iff the current node has just received a response for its type- i request.

Consider the following *global* properties, where N is the set of possible request types:

⁶<https://github.com/Harniver/past-ctl-monitoring>

1. *timeout*: every request of type i should be followed by a suitable response (i.e., of type i) from some other node within n time instants. In a negative form, we can express it as:

$$\neg \text{EP}(\text{no_reply}(rq_i, rs_i, n))$$

where *no_reply* is defined as:⁷

$$\text{no_reply}(rq, rs, n) = \begin{cases} rq & \text{if } n = 0 \\ \neg rs \wedge Y \text{no_reply}(rq, rs, n - 1) & \text{otherwise.} \end{cases}$$

Note that we are actually defining as many properties as the number of request types N .

2. *spurious*: a response of any type $i \in N$ should be received only if a corresponding request was issued at some previous time instant. In formulas:

$$\bigwedge_{i \in N} \text{AH}(rs_i \Rightarrow Y(\neg rs_i S rq_i))$$

505 This formula has already been considered in Examples 3.1 - 3.3. Notice that here we consider a single formula for all the types $i \in N$ (external \wedge ranging on $i \in N$).

3. *double req*: nodes issue requests one at a time, i.e., they wait for a response before issuing another request. Again in a negative form, and considering a single formula for all $i \in N$, we can express it as:

$$\bigwedge_{i \in N} \neg \text{EP}(Y((\neg rs_i S rq_i) \wedge rq_i))$$

We evaluated this scenario in a logical environment, depicted in a 2D plane for presentation purposes (see Figure 8 for a snapshot), where network nodes are as follows: 5 central cloud nodes, all connected to each other; 20 intermediate fog nodes, each connected with 2 or 3 cloud nodes; and 50 edge nodes each
510 connected with 3 or 4 fog nodes. Note that neither fog nodes, nor edge nodes can communicate with their peers. The network topology of cloud and fog nodes is static, while every edge node appears and disappears at a random time after the beginning of the simulation.

Nodes move back and forth between two states: *computing* (white nodes) and *pending request* (coloured nodes). The transition from *computing* to *pending request* happens when the node sends a request rq_i ; the
515 opposite transition happens when the node receives a response rs_i .

The type of the current request is represented by the colour of the node: yellow, green, red, and fuchsia for types 1 to 4 (white nodes do not have a pending request). As the node spends time waiting for a response, its sides become darker. The simulator has different (average) response times for the different request types, with the time for type 1 being the highest, and the others progressively lower. As expected, the monitors of

⁷It should be clear that the recursive definition of *no_reply* is easily expressed in past-CTL once n has been fixed.




Figure 8: Screenshot of a simulated network. Coloured nodes have a pending request (colour based on request type, darker sides for older requests), square nodes if the monitor is failed.

520 the *timeout* properties detect a timeout for type 1 earlier than for types 2 and 3, while no timeout for type 4 is detected in the simulation horizon; this is illustrated by Figure 9, which plots, for each request type, the rate of nodes that know that the *timeout* property has been violated. Note that, in Figure 8, the shape of nodes that know about a timeout violation changes from sphere to cube.

In order to test the monitoring of the *spurious* property above, we have injected the possibility for a
 525 node to receive a reply not matching the type of its last request. In particular, the solid black line in the graph of Figure 9 shows that the violation of the *spurious* property is quickly propagated to all the nodes after time 40.

Similarly, for testing the monitoring of the *double req* property, we have injected the possibility for a
 530 node to occasionally issue a second request before receiving a response. The dashed orange line in the graph of Figure 9 shows that the violation of the *double req* property is quickly propagated to all the nodes at a time close to 100.

4.2. Drones Recognition

Suppose that N areas need to be handled by drones in some way (perform a recognition, spray substances, and similar). Proposition q_i for $i = 1 \dots N$ is true iff the current drone is now handling area i . The properties
 535 below are formulated in terms of those observables, and how drones actually move or communicate are external factors for the purpose of our example. The existential quantifiers also neatly encode that at the

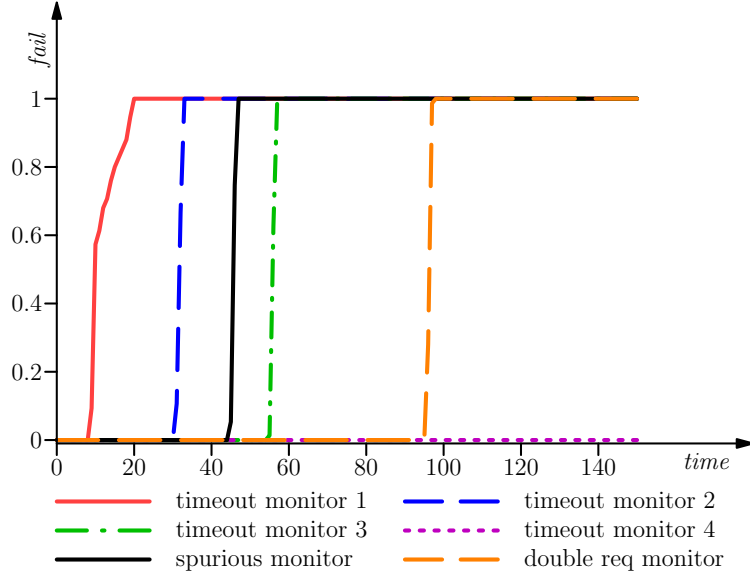


Figure 9: Monitors evolution in time.

same point in time, different drones may have different views of the global system. Consider the properties:

- Every area gets eventually handled by a drone (true means success, false means “not yet”): $\bigwedge_{i \in N} \text{EP } q_i$.
Observe the requirements that this property imposes on the underlying event structure: if e.g. we would like to use this property as a criterion for distributed termination, we have to make sure that we implement a system where eventually every drone will somehow see the past of every other drone.
- No drone is handling an area that it knows to be already handled (true means “ok so far”, false means failure). In other words, it never happens that a drone is handling an area that it knows to be already handled: $\bigwedge_{i \in N} \text{AH } \neg(q_i \wedge \text{EY}(\text{EP } q_i))$.

The first property can be understood as *liveness*, while the second as *safety*.

We evaluated this scenario in an open 3D environment $1000m \times 1000m \times 100m$ large (see Figure 10 (left) for a snapshot), subdivided into four square areas each with a communication tower at its middle (brown/red cubes). 50 drones swept the area to perform random tasks (green cubes are drones heading to a destination, yellow ones are handling their goal), interleaved by recharging (blue) and waiting for tasks (grey) periods. We assumed that drones and towers followed an edge-like connection topology: towers were all connected with each other, and drones were connected to the one or two towers that were close enough to them.

The four towers issued “handling”-requests at a random time between 0s and 200s, instructing the closest waiting drones to reach them. Shortly after having issued the request, every tower’s request was served and the corresponding handling monitor (liveness) quickly started becoming true in every node in the network (see Figure 10 (right), showing the evolution in time of the percentage of nodes evaluating the monitors to

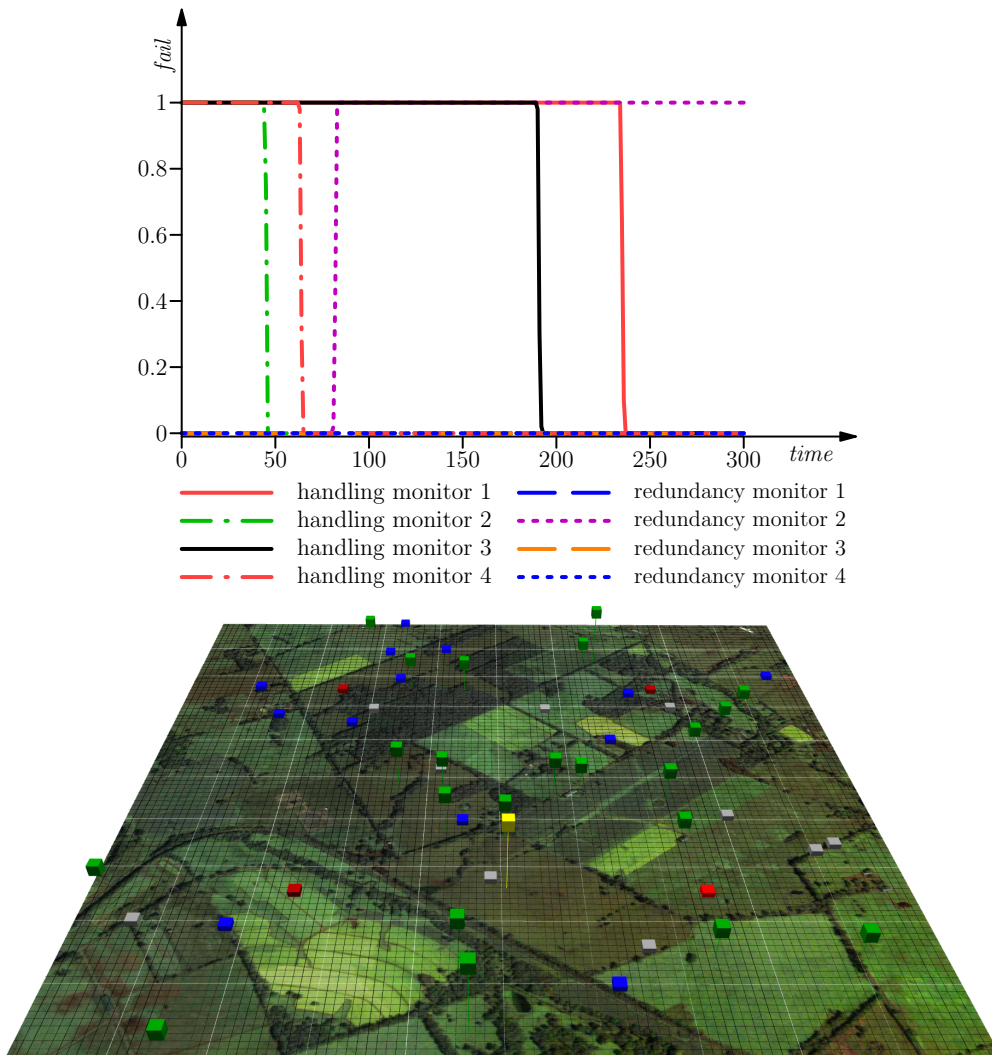


Figure 10: Screenshot of a simulated drone area (bottom), monitors evolution in time (top). Brown cubes are communication towers (red if asking service), others are drones (green-moving, yellow-handling, blue-recharging, grey-idling).

\perp). Due to a fault in the algorithm, a second drone reached and handled tower 2 at $t = 80$, causing the corresponding safety monitor to quickly become false in every device of the network. The safety monitors of the other three areas remained true for the whole simulation time, since no other re-servicing occurred.

560 Notice that an individual breakdown of the propositions composing the monitor through boolean operators allows for an easier understanding of the precise behaviour of the system. This is particularly true if a proposition is composed from a liveness and safety property together: the conjunction of such properties may be continuously false, and we could not tell apart situations in which liveness has not been reached yet, from situations in which liveness has been reached only while breaking safety requirements.

Previous works [5] present field-calculus solutions to the problem of safe crowd dispersal for disaster management. Armed with our new formalism, we can now specify the property that every agent does not enter back into the unsafe area after having evacuated from the unsafe area. The all-quantification here is implicit, as we assume that every agent is evaluating the same property, and only considering its own past. For this example, we assume a sensor (or field-calculus function) *safe* in every agent, and an indicator for whether an evacuation *alert* has been given, and the device should suggest to move from an unsafe area to safety. We specify that whenever in safety during an alert, we do not regress. As such, we do not use any branching-time features, and we can specify the property as:

$$\varphi_1 := \mathbf{H}(\mathbf{Y}(\textit{safe} \wedge \textit{alert}) \rightarrow (\textit{safe} \vee \neg \textit{alert})).$$

It is easy to see how the truth value of this property may evolve. Starting out in an unsafe location, the overall verdict remains *true* even once safety is reached. It is also obvious that the verdict may switch to *false* depending on external factors, such as the agent moving or the unsafe area encroaching the agent.

In a next step, we could then lift that monitor to an observer which specifies that everyone in the past cone of events has reached safety without intermittently getting into danger again. We achieve this by switching out the linear-time past operator \mathbf{H} to the branching-operator \mathbf{AH} :

$$\varphi_2 := \mathbf{AH}(\mathbf{Y}(\textit{safe} \wedge \textit{alert}) \rightarrow (\textit{safe} \vee \neg \textit{alert})).$$

Note that the verdict *true* then may become invalidated for agents that learn from/come in contact with agents which themselves observed violations of this property over time.

We evaluated this scenario in a simulated $500m \times 500m$ area (see Figure 11 (left) for a screenshot). We randomly choose 5 potential incident points (black cubes, that switch to red during an incident) and 100 people randomly walking through the area; the colour of the cubes representing people varies from green to blue (when they are *safe*) and from red to yellow (when they are *not safe*), depending on the distance from an incident. A person is safe when she is more than $50m$ from the incident. The distance from the incident is not directly sensed by the people, but is computed through a distributed Field Calculus function that only assumes knowledge of the distance of a person from her neighbours. Units on people were able to communicate within a $50m$ range, while incident points had an higher communication range ($100m$ with people, $200m$ among them).

During the simulation, an incident happens in each of the 5 points, at a random time from 0 to $300s$, and lasts for a random duration from 0 to $100s$. Note that in this way incidents can, and in fact often do in practice, overlap with each other. Finally, an *alert* is perceived by people when they are within a $100m$ radius from an incident. The reaction of a person being *unsafe* is to immediately start moving towards her neighbour that is further away from the incident, until she is at least $10m$ outside the unsafe area, i.e., $60m$

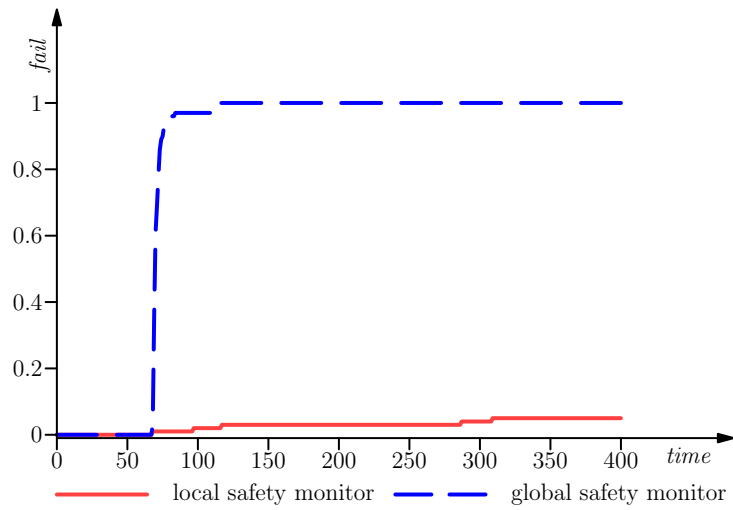


Figure 11: Screenshot of a simulated incident area (bottom), monitors evolution in time (top). Incident points are black cubes (red during incidents), peoples' colour varies from green to blue (if safe) and from red to yellow (if unsafe) depending on distance. Cubes are larger where the monitor has failed.

585 from the incident. Then, she goes back to random walk, which is an important factor for making it possible to violate local safety.

After the first incident at about 60s, few people violate formula φ_1 of local safety. However, such a violation quickly spreads to most other people so that the indicator of *global* safety failure almost immediately reaches 100% (Figure 11 (right), blue dashed line); this means that most or all of the people know that someone has behaved unsafely, i.e., entered the unsafe area after being in a safe place. The *local* safety failure indicator only increases when, during the five incidents, specific persons behave unsafely, and involves less than 10% of the people by the end of the simulation. Note that the introduction of the *all-paths* quantifier in this example creates the impression of a global observer, yet we have to recall that all preceding branches 590 are derived from the event structure, and not every agent may be in contact with everyone else.

Finally, we consider a smart home scenario borrowed from a preliminary work on using field calculus for runtime verification [3, Sect. 4.2], and hypothetically situated in a large building where multiple sensors are connected through a multi-hop network. Assume that we want to check that an electronic system (e.g. a room light) is active whenever some people have been present in the immediate vicinity in the close past.⁸

600 Consider the predicates:

- s is true if the current device is responsible of an electronic system;
- a is true if s is true and the system is active;
- p is true if the current device is detecting the presence of people.

The required property may be expressed as $\text{AH } \phi_s$ where $\phi_s = s \rightarrow (p \leftrightarrow a)$, however, that would not accommodate for any (even small) response delay for the system. If we want to ensure that a is true iff p is true now and in the previous round (resp. with false), then we can change the property to $\text{AH } \phi_w$ where:

$$\phi_w = s \rightarrow (p \wedge \text{Y } p \rightarrow a) \wedge (\neg p \wedge \text{Y } \neg p \rightarrow \neg a).$$

We evaluated this scenario in a simulated building floor comprising eight $6m \times 6m$ rooms, each with a central light, around a $24m \times 3m$ corridor with four lights (see Figure 12 (left) for a screenshot). We populated 605 the building with 12 people (pink cubes), randomly walking to a different room through the corridor after stops of 40s in average (according to an exponential distribution), and assumed a short connection range of 3m modelling a low-power Bluetooth system. Lights turned on (yellow) or off (grey) according to the presence of people within the connection range (predicate p), with a 5% chance of a delayed reaction during an on/off switch, and a lower 0.3% spurious reaction during rounds in which p was steady. Due to these 610 error probabilities, the local strong (ϕ_s) and weak (ϕ_w) monitors became occasionally false on individual lights during the simulated time (small spikes in the graph). After the first error signalled by ϕ_s at about 1s in the simulation, the global monitor $\text{AH } \phi_s$ started becoming false; with the spreading of its value across the network being limited by the low connectivity of the network. A similar situation happened with ϕ_w 615 and $\text{AH } \phi_w$ at about 27s. The screenshot represents an intermediate state, where some devices are aware of the failure (larger cubes) while other still are not yet.

5. Related Work

5.1. Aggregate Computing

The problem of finding suitable programming models for ensemble of devices has been the subject of 620 intensive research—see e.g. the surveys [19, 15]): works as TOTA [20] and Hood [21] provide abstractions

⁸In [3], the property considered a single light active whenever people were present somewhere in the building.

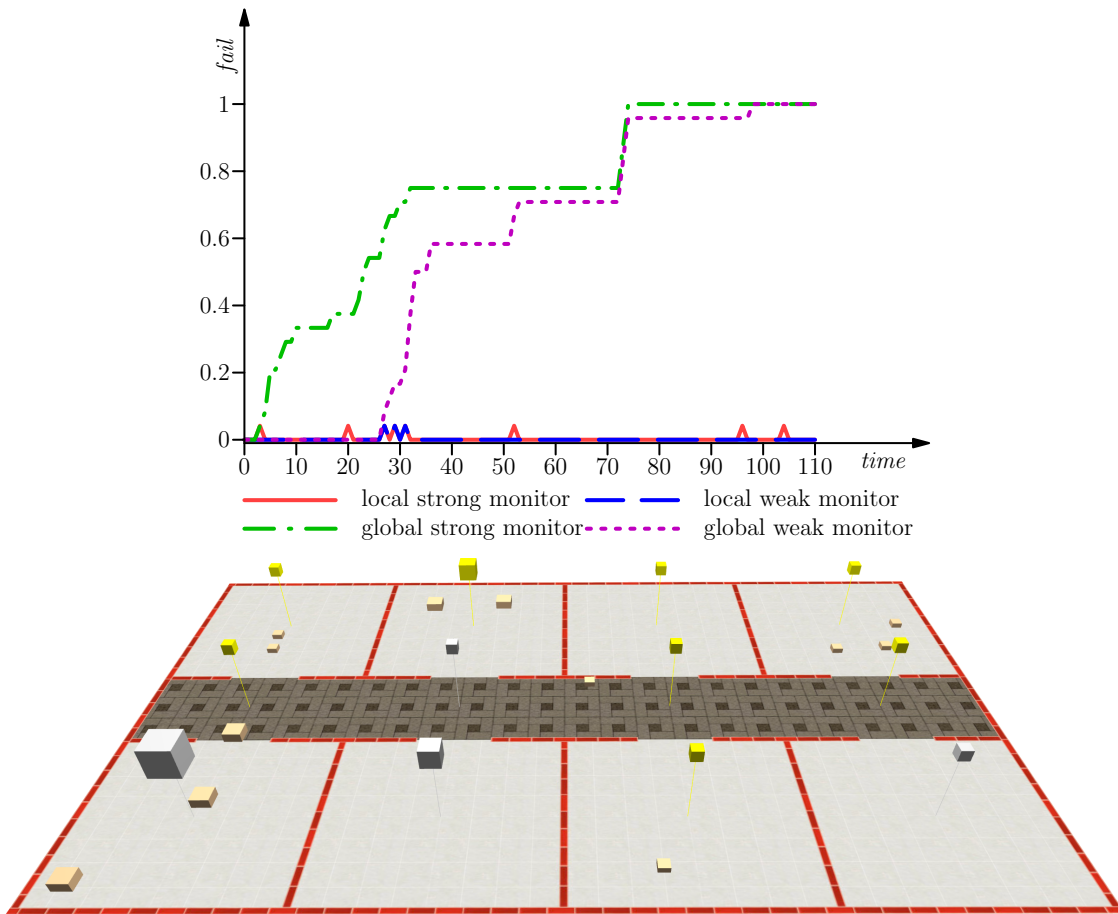


Figure 12: Screenshot of a simulated building floor (bottom), monitors evolution in time (top). Lights are higher cubes (grey if off, yellow if on), lower cubes are people (larger cubes where the monitor has failed).

over the single device to facilitate construction of macro-level systems; GPL [22] and others are used to express spatial and geometric patterns; Regiment [23] is an information system used to stream and summarise information over space-time regions; while MGS [24] and the fixpoint approach in [25] provide general purpose space-time computing models. Aggregate computing and the field calculus have then been developed as a generalisation of the above approaches, with the goal of defining a programming model with sufficient expressiveness to describe complex distributed processes by a functional-oriented compositional model, whose semantics is defined in terms of gossip-like computational processes.

Hence, *aggregate computing* [5] aims at supporting reusability and composability of collective adaptive behaviour as inherent properties. Following the inspiration of “fields” of physics (e.g., gravitational fields), this is achieved by the notion of *computational field* (simply called *field*) [20], defined as a global data structure mapping devices of the distributed systems to computational values. Computing with fields means deriving in a computable way an output field from a set of input fields. This can be done at a low-level, to defined programming language constructs or general-purpose building blocks or reusable behaviour, or at

a high-level to design collective adaptive services or whole distributed applications—which ultimately work
635 by getting input fields from sensors and process them to produce output fields to actuators.

The *field calculus* [6] is a minimal functional language that identifies basic constructs to manipulate fields, and whose operational semantics can act as blueprint for developing toolchains to design and deploy systems of possibly myriad devices interacting via proximity-based broadcasts—see e.g. the FCPP, PROTELIS and SCAFI DSLs. The field calculus have been used to formally prove distributed systems properties of resiliency
640 to environment changes [26, 27, 13] and to device distribution [28], which makes it particular suitable for distributed runtime verification.

5.2. Distributed Runtime Verification

Runtime verification is a lightweight verification technique concerned with observing the execution of a system with respect to a specification [1]. Specifications are generally trace- or stream-based, with events
645 that are mapped to atomic propositions in the underlying logic of the specification language. Popular specification languages include variations on the Linear Time Logic (LTL), and regular expressions, which can be effectively checked through finite automata constructions. Events may be generated through state changes or execution flow, such as method calls.

In *distributed* runtime verification, this concept is lifted to distributed systems, to meet applications in
650 areas like: (i) observing distributed computations and expressiveness (specifications over the distributed systems), (ii) analysis decomposition (coupled composition of system- and monitoring components), (iii) exploiting parallelism (in the evaluation of monitors), (iv) fault tolerance and (v) efficiency gains (by optimising communication) [2].

Naturally, such lifting also affects the specification language. Bauer and Falcone [29] show a decentralised
655 monitoring approach where disjoint atomic propositions in a global LTL property are monitored without a central observer in their respective components. Communication overhead is shown to be lower than the number of messages that would need to be sent to a central observer.

Sen *et al.* introduce PT-DTL [30] to specify distributed properties in a past time temporal logic. Sub-
formulas in a specification are explicitly annotated with the node (or process) where the subformula should
660 be evaluated. Communication of results of subcomputation is handled by message passing. A future time logic with explicit annotations (and a three-valued logic) based on the same idea has been presented in [31].

The above approaches assume a total communication topology, i.e., each node can send messages to everyone in the system, although causally unrelated messages may arrive in arbitrary order. Given the dynamic nature of IoT and edge scenarios, such a fixed distribution of verification sub-tasks is not possible
665 in our approach, and the hierarchical structure does not further our aim of making each node an instance of the monitor.

As any other software engineering technique, also runtime verification in general has been applied in some form to IoT and edge computing. Tsigkanos *et al.* [32] conduct a case-study in monitoring edge-based systems with temporal MFOTL specification in a similar explicit hierarchical setup. Unlike our approach, 670 their network configuration is static and each tier evaluates its own specification, and specifications explicitly refer to subformulas on other tiers. The latter distinction is only minor: our single field calculus specification would guard subformulas for the tier they apply to, and at the same time solve the problem of keeping signatures across formulas in sync.

Inçki *et al.* [33] specify expected behaviour in an IoT system through message sequence charts and then 675 derive a monitor over traces that is centrally checked with a complex event processing tool. Unlike our solution, their approach relies on having network sniffers observing messages in each node and sending events to the central monitor.

We position ourselves in the following in the design space of distributed and decentralised runtime verification (see Francalanza *et al.* [2]): we assume that every agent, as a constituent of the whole system, 680 executes independently and occasionally synchronises or communicates with other agents via the underlying communication platform. Going with the analogy of taking an agent as a *process* in Francalanza’s terminology, we consider any two processes as *remote* to each other. A *local trace of events* corresponds to a sequence of sets of values for observables, as defined through the sensors of an agent, or derived values from those. Since agents may appear or disappear over time from the overall system, traces from different processes 685 are not aligned in time in the sense that for a particular index/position in each trace, these events did not necessarily happen at the same time. Although the event structure may give the impression of a post-hoc view, the evaluation strategy that we have proposed is decidedly online: the field calculus expressions that evaluate branching expressions rely on communication with neighbouring agents (at that point in time).

Monitoring is performed by computation entities, identified by *monitors* M , that check properties of the 690 system under analysis by analysing the traces. Similar to processes each monitor is hosted at a given location and may communicate with other monitors. As every agents is executing the same field calculus program, and hence the same monitor, our proposed approach brings a conceptual simplification. Namely, despite the distributed setting it resembles more the traditional setup of runtime verification: from the perspective of the programmer there is a single monitor and a single trace is evaluated—though this monitor is deployed 695 and executed on a network of distributed devices and this trace contain events from different devices.

Moreover, being based on the field calculus, our approach is intrinsically resilient—cf. the discussion at the end of Sect. 5.1. Namely, it self-adapts to changes in the network topology. So, it automatically handles *failures*, which usually make distributed systems harder to manage: a non-responsive node (through a crash or through intermittently unavailable communication) only affects others through its lack in participation 700 in evaluation of field expressions, which may affect the verdicts computed by its neighbours. We do not explicitly consider faulty sensors or message corruption, but leave this to integrity measures e.g. on the

communication layer.

The absence of a global verdict, and instead having a per-node instantiation also puts us at odds with respect to another recent classification effort of distributed and decentralised runtime verification[34, 35].

705 Such work considers challenges and approaches to monitor decomposition, placement and control (which does not apply here), and fault tolerance, for which we have made the case in the introduction that our approach presents a natural solution to. Many of the known issues in distributed runtime verification do not arise in our approach, due to the conceptual restriction of having to run the same monitor on every node.

6. Conclusion and Future Work

710 We provided a compositional translation of properties expressed in past-CTL logic into field calculus programs for monitoring them, proving that the translated monitor runs using local memory, message size and computation time that are all linear in the size of the formula (1 bit per temporal connective), matching the efficiency of previous result for past-LTL in a non-distributed setting [8]. We then provided sample applications from different scenarios of edge computing (disaster management, swarm control, smart home),
715 validating the monitor effectiveness by simulation under faults and changes, and arguing that past-CTL logic is reasonably expressive for properties of distributed systems.

In future work, we plan to extend this evaluation to scenarios based on real GPS traces, possibly expanding existing case studies in built environments [36], also considering the use of past-CTL properties in conjunction with runtime *reflection* [1], where a property is annotated with a desired action (such as activation of an actuator) when it is activated, in order to manipulate the behaviour of agents for controlling
720 the overall system evolution.

Furthermore, parallel studies have been investigating the usage of field calculus for monitoring SLCS spatial logic formulas [37]. Despite aiming towards a similar purpose, this work shares very little with the present one: the two logics have no modal operators in common, and the models for interpreting them are
725 incompatible. While past-CTL formulas are interpreted in event structures, SLCS formulas are interpreted in *graphs* (sets of nodes and vertices) without a temporal component, which are intended to represent the status of an evolving network in a particular instant. Even though time is abstracted away in the model of SLCS formulas, it still plays an implicit role, since the computation of the value of such formulas requires some time to converge and adjust. Expanding on a recent preliminary discussion on this topic [38], in future
730 work we plan to formalise the impact of time on SLCS formulas, giving an interpretation of them in event structures and thus allowing its merger with past-CTL, obtaining a full-fledged spatio-temporal logic.

Finally, we plan to investigate whether a distributed *continuation* operator could be devised to allow the monitoring of a future LTL-like logic in field calculus, in order to address settings where stronger prediction capabilities are needed and resources are not tightly constrained.

735 **References**

- [1] M. Leucker, C. Schallhart, A brief account of runtime verification, *J. Log. Algebr. Program.* 78 (5) (2009) 293–303 (2009). doi:10.1016/j.jlap.2008.08.004.
- [2] A. Francalanza, J. A. Pérez, C. Sánchez, Runtime verification for decentralised and distributed systems, in: *Lectures on Runtime Verification - Introductory and Advanced Topics*, Vol. 10457 of *Lecture Notes in Computer Science*, Springer, 2018, pp. 176–210 (2018). doi:10.1007/978-3-319-75632-5_6.
- 740 [3] G. Audrito, F. Damiani, V. Stolz, M. Viroli, On distributed runtime verification by aggregate computing, in: *2nd Workshop on Verification of Objects at RunTime EXecution (VORTEX)*, Vol. 302 of *EPTCS*, 2018, pp. 47–61 (2018). doi:10.4204/EPTCS.302.4.
- [4] W. Shi, J. Cao, Q. Zhang, Y. Li, L. Xu, Edge computing: Vision and challenges, *IEEE Internet Things J.* 3 (5) (2016) 637–646 (2016).
- 745 [5] J. Beal, D. Pianini, M. Viroli, Aggregate programming for the Internet of Things, *IEEE Computer* 48 (9) (2015) 22–30 (2015). doi:10.1109/MC.2015.261.
- [6] G. Audrito, M. Viroli, F. Damiani, D. Pianini, J. Beal, A higher-order calculus of computational fields, *ACM Trans. Comput. Log.* 20 (1) (2019) 5:1–5:55 (2019). doi:10.1145/3285956.
- 750 [7] S. Taherizadeh, A. C. Jones, I. Taylor, Z. Zhao, V. Stankovski, Monitoring self-adaptive applications within edge computing frameworks: A state-of-the-art review, *J. Systems and Software* 136 (2018).
- [8] K. Havelund, G. Rosu, Synthesizing monitors for safety properties, in: *Tools and Algorithms for the Construction and Analysis of Systems (TACAS)*, Vol. 2280 of *Lecture Notes in Computer Science*, Springer, 2002, pp. 342–356 (2002). doi:10.1007/3-540-46002-0_24.
- 755 [9] G. Audrito, FCPP: an efficient and extensible field calculus framework, in: *International Conference on Autonomic Computing and Self-Organizing Systems (ACSOS)*, IEEE, 2020, pp. 153–159 (2020). doi:10.1109/ACSOS49614.2020.00037.
- [10] D. Pianini, M. Viroli, J. Beal, Protelis: practical aggregate programming, in: *Proceedings of the 30th Annual ACM Symposium on Applied Computing*, ACM, 2015, pp. 1846–1853 (2015). doi:10.1145/2695664.2695913.
- 760 [11] R. Casadei, M. Viroli, G. Audrito, F. Damiani, Fscafi : A core calculus for collective adaptive systems programming, in: *Leveraging Applications of Formal Methods, Verification and Validation: Engineering Principles*, Vol. 12477 of *Lecture Notes in Computer Science*, Springer, 2020, pp. 344–360 (2020). doi:10.1007/978-3-030-61470-6_21.
- [12] G. Audrito, J. Beal, F. Damiani, M. Viroli, Space-time universality of field calculus, in: *Coordination Models and Languages*, Vol. 10852 of *Lecture Notes in Computer Science*, Springer, 2018, pp. 1–20 (2018). doi:10.1007/978-3-319-92408-3_1.
- 765 [13] M. Viroli, G. Audrito, J. Beal, F. Damiani, D. Pianini, Engineering resilient collective adaptive systems by self-stabilisation, *ACM Transactions on Modelling and Computer Simulation* 28 (2) (2018) 16:1–16:28 (2018). doi:10.1145/3177774.
- [14] L. Lamport, Time, clocks, and the ordering of events in a distributed system, *Commun. ACM* 21 (7) (1978) 558–565 (1978). doi:10.1145/359545.359563.
- 770 [15] M. Viroli, J. Beal, F. Damiani, G. Audrito, R. Casadei, D. Pianini, From distributed coordination to field calculus and aggregate computing, Vol. 109, 2019 (2019). doi:10.1016/j.jlamp.2019.100486.
- [16] G. Audrito, J. Beal, F. Damiani, D. Pianini, M. Viroli, Field-based coordination with the share operator, *Log. Methods Comput. Sci.* 16 (4) (2020). doi:10.23638/LMCS-16(4:1)2020.
- [17] N. Gigante, A. Montanari, M. Reynolds, A one-pass tree-shaped tableau for LTL+past, in: *Logic for Programming, Artificial Intelligence and Reasoning (LPAR)*, Vol. 46 of *EPiC Series in Computing*, EasyChair, 2017, pp. 456–473 (2017). URL <http://www.easychair.org/publications/paper/340363>

- [18] R. V. Levine, A. Norenzayan, The pace of life in 31 countries, *Journal of cross-cultural psychology* 30 (2) (1999) 178–205 (1999).
- [19] J. Beal, S. Dulman, K. Usbeck, M. Viroli, N. Correll, Organizing the aggregate: Languages for spatial computing, in: *Formal and Practical Aspects of Domain-Specific Languages: Recent Developments*, IGI Global, 2013, Ch. 16, pp. 436–501 (2013). doi:10.4018/978-1-4666-2092-6.ch016.
- [20] M. Mamei, F. Zambonelli, Programming pervasive and mobile computing applications: The TOTA approach, *ACM Trans. on Software Engineering Methodologies* 18 (4) (2009) 1–56 (2009). doi:10.1145/1538942.1538945.
- [21] K. Whitehouse, C. Sharp, D. E. Culler, E. A. Brewer, Hood: A neighborhood abstraction for sensor networks, in: G. S. Banavar, W. Zwaenepoel, D. Terry, R. Want (Eds.), *2nd International Conference on Mobile Systems, Applications, and Services*, ACM / USENIX, 2004 (2004). doi:10.1145/990064.990079.
- [22] D. Coore, *Botanical computing: A developmental approach to generating inter connect topologies on an amorphous computer*, Ph.D. thesis, MIT, Cambridge, MA, USA (1999).
- [23] R. Newton, M. Welsh, Region streams: Functional macroprogramming for sensor networks, in: *Workshop on Data Management for Sensor Networks*, 2004, pp. 78–87 (2004). doi:10.1145/1052199.1052213.
- [24] J. Giavitto, O. Michel, J. Cohen, A. Spicher, Computations in space and space in computations, in: *Unconventional Programming Paradigms*, Vol. 3566 of *Lecture Notes in Computer Science*, Springer, 2004, pp. 137–152 (2004). doi:10.1007/11527800_11.
- [25] A. Lluch-Lafuente, M. Loreti, U. Montanari, Asynchronous distributed execution of fixpoint-based computational fields, *Logical Methods in Computer Science* 13 (1) (2017). doi:10.23638/LMCS-13(1:13)2017.
- [26] G. Audrito, F. Damiani, M. Viroli, E. Bini, Distributed real-time shortest-paths computations with the field calculus, in: *IEEE Real-Time Systems Symposium (RTSS)*, IEEE Computer Society, 2018, pp. 23–34 (2018). doi:10.1109/RTSS.2018.00013.
- [27] Y. Nishiwaki, F-calculus: A universal programming language of self-stabilizing computational fields, in: *1st Intl. Workshops on Foundations and Applications of Self* Systems (FAS*W)*, IEEE, 2016, pp. 198–203 (2016). doi:10.1109/FAS-W.2016.51.
- [28] J. Beal, M. Viroli, D. Pianini, F. Damiani, Self-adaptation to device distribution in the Internet of Things, *ACM Transaction on Autonomous and Adaptive Systems* 12 (3) (2017) 12:1–12:29 (2017). doi:10.1145/3105758.
- [29] A. Bauer, Y. Falcone, Decentralised LTL monitoring, *Formal Methods in System Design* 48 (1-2) (2016) 46–93 (2016). doi:10.1007/s10703-016-0253-8.
- [30] K. Sen, A. Vardhan, G. Agha, G. Rosu, Efficient decentralized monitoring of safety in distributed systems, in: *26th Intl. Conf. on Software Engineering*, 2004, pp. 418–427 (2004). doi:10.1109/ICSE.2004.1317464.
- [31] T. Scheffel, M. Schmitz, Three-valued asynchronous distributed runtime verification, in: *Twelfth ACM/IEEE Conference on Formal Methods and Models for Codesign (MEMOCODE)*, 2014 (2014).
- [32] C. Tsigkanos, M. M. Bersani, P. A. Frangoudis, S. Dustdar, Edge-based runtime verification for the Internet of Things, *IEEE transactions on services computing* (2021) 1–1 (2021).
- [33] K. Inçki, I. Ari, A novel runtime verification solution for IoT systems, *IEEE Access* 6 (2018) 13501–13512 (2018). doi:10.1109/ACCESS.2018.2813887.
- [34] C. Sánchez, G. Schneider, W. Ahrendt, E. Bartocci, D. Bianculli, C. Colombo, Y. Falcone, A. Francalanza, S. Krstic, J. M. Lourenço, D. Nickovic, G. J. Pace, J. Rufino, J. Signoles, D. Traytel, A. Weiss, A survey of challenges for runtime verification from advanced application domains (beyond software), *Formal Methods Syst. Des.* 54 (3) (2019) 279–335 (2019). doi:10.1007/s10703-019-00337-w.
- [35] A. El-Hokayem, Y. Falcone, On the monitoring of decentralized specifications: Semantics, properties, analysis, and simulation, *ACM Trans. Softw. Eng. Methodol.* 29 (1) (Jan. 2020). doi:10.1145/3355181.

- 820 [36] G. Audrito, F. Damiani, G. M. D. Giuda, S. Meschini, L. Pellegrini, E. Seghezzi, L. C. Tagliabue, L. Testa, G. Torta, RM for users' safety and security in the built environment, in: Proceedings of the 5th ACM International Workshop on Verification and mOnitoring at Runtime EXecution (VORTEX), ACM, 2021, pp. 13–16 (2021). doi:10.1145/3464974.3468445.
- [37] G. Audrito, R. Casadei, F. Damiani, V. Stolz, M. Viroli, Adaptive distributed monitors of spatial properties for cyber-physical systems, *Journal of Systems and Software* 175 (2021) 110908 (2021). doi:j.jss.2021.110908.
- 825 [38] G. Audrito, G. Torta, Towards aggregate monitoring of spatio-temporal properties, in: Proceedings of the 5th ACM International Workshop on Verification and mOnitoring at Runtime EXecution (VORTEX), ACM, 2021, pp. 26–29 (2021). doi:10.1145/3464974.3468448.