

On Distributed Runtime Verification by Aggregate Computing

Giorgio Audrito

Ferruccio Damiani

University of Turin, Italy

giorgio.audrito@gmail.com

ferruccio.damiani@unito.it

Volker Stolz

Mirko Viroli

Western Norway University of Applied Sciences, Norway

University of Bologna, Italy

Volker.Stolz@hvl.no

mirko.viroli@unibo.it

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. Distributed runtime verification is runtime verification in connection with distributed systems: it comprises both monitoring of distributed systems and using distributed systems for monitoring. Aggregate computing is a programming paradigm based on a reference computing machine that is the aggregate collection of devices that cooperatively carry out a computational process: the details of behaviour, position and number of devices are largely abstracted away, to be replaced with a space-filling computational environment. In this position paper we argue, by means of simple examples, that aggregate computing is particularly well suited for implementing distributed monitors. Our aim is to foster further research on how to generate aggregate computing monitors from suitable formal specifications.

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. Distributed runtime verification is runtime verification in connection with distributed systems: it comprises both monitoring of distributed systems and using 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. 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), 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, though the property to verify is global). 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 argue that a promising approach to address these challenges can be rooted on the computational paradigm of aggregate computing [11], along with the field calculus language [17].

Aggregate computing promotes a view of distributed systems as a conceptually single computing device, spread throughout the (physical or virtual) space in which nodes are deployed. At the paradigm level, hence, this view promotes the specification (construction, reasoning, programming) of global-level computational behaviour, where the interaction of individuals are essentially abstracted away. At the modelling level, the field calculus can be leveraged, which expresses computations as transformations of *computational fields* (or *fields* of short), namely, space-time distributed data structures mapping computational events (occurring at a given position of space and time) to computational values. 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. Field calculus has a working implementation called SCAFI in the Scala programming language [14], where field computations can be expressed by Scala functions (relying on a suitable API) and actors are generated to realise the distributed system.

The remainder of this paper is organized as follows: Section 2 provides the necessary background; Section 3 presents the field calculus; Section 4 discusses monitoring general distributed programs through field calculus; Section 5 illustrates how field calculus programs can be instrumented with monitors; and Section 6 concludes.

2 Background

In this section we provide the necessary background on distributed runtime verification and aggregate computing by briefly discussing the literature outlining their role.

2.1 Motivating examples

We motivate the usage of aggregate computing together with runtime verification techniques through two examples that have been thoroughly studied in earlier literature: a crowd-evacuation scenario and a general communication channel.

In the first scenario, a program (not necessarily written through aggregate computing techniques) is used to manage evacuation of agents from a given area in case of an emergency. Ideally, in such critical situations correctness guarantees for a particular solution and its implementation would be needed. Since the guarantees that can be proved are usually not fully satisfactory, they can be fruitfully complemented with runtime monitors. As an example, we focus here on a simple “per-agent” property that we could monitor: *two neighbour agents* (closer than 5m) *should not have “evacuation vectors” that lead to a direct collision with each other* (for both agents, the evacuation vector is within 60° from the direction of the other agent). We can then instantiate an aggregate computing monitor on each agent, observing the local and neighbours’ evacuation vectors, and flagging violations of this property as they occur.

The second scenario describes an aggregate computing solution to the well-known problem of establishing a shortest communication path between two nodes, while ensuring reliability through an imposed *width* (cross-section size of the channel), which provides the desired redundancy and alternative routes. In this situation, we define a more interesting property that is not per-agent, but rather *per-network*, and also shows how the monitor can feedback into the original program: for each cross-section, we require that it has at least *min* width of alternative connections. If not, we demand the channel program to increase the width. Conversely, if all slices contain more than *max* nodes, we shrink the channel to save computational power.

2.2 Distributed runtime verification

Runtime verification is a lightweight verification technique concerned with observing the execution of a system with respect to a specification [24]. Specifications are generally trace- or stream-based, with events that are mapped to atomic propositions in the underlying logic of the specification language. Popular specification languages include variations on the Linear Temporal 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, we lift this concept to distributed systems, where we find applications in the following areas [20]: (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). In Sections 4 and 5, we show how runtime verification can be applied to, or contribute to some of those areas.

Naturally, such lifting also affects the specification language. Bauer and Falcone [9] show a decentralised 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 [29] 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 sub-formula should be evaluated. Communication of results of sub-computation is handled by message passing.

Both 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.

Going beyond linear-time properties, hyperproperties over a set of traces allow a richer expressivity [19]. In our setting, as each node is running the same program, we can understand such a set as consisting of traces from the individual nodes. Further issues on (efficient) monitorability have been addressed by Aceto *et al.* in [1].

2.3 Aggregate computing

The problem of finding suitable programming models for ensemble of devices has been the subject of intensive research—see e.g. the surveys [10, 32]: works as TOTA [26] and Hood [33] provide abstractions over the single device to facilitate construction of macro-level systems; GPL [15] and others are used to express spatial and geometric patterns; Regiment [27] and TinyLime [16] are information systems used to stream and summarise information over space-time regions; while MGS [21] and the fixpoint approach in [25] provide general purpose space-time computing models. Aggregate computing and the field calculus have then be 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* [11] 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*) [26], defined as a global data structure mapping devices of the distributed system 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, by defining programming language constructs or general-purpose building blocks

of reusable behaviour, or at a high-level by designing collective adaptive services or whole distributed applications—which ultimately work by getting input fields from sensors and process them to produce output fields to actuators.

The *field calculus* [17, 31] 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. Recent works have also adopted this field calculus as a *lingua franca* to investigate formal properties of resiliency to environment changes [28, 31], and to device distribution [12].

2.4 Deployment

A number of techniques exist to deploy runtime verification as part of or in parallel to an application to be subjected to runtime verification. A high-level technique to monitor a JVM-based application is the use of aspect-oriented programming [30], which allows for an easy integration in terms of events: this method allows to easily intercept actions of the main application and use them as input events for the step-wise evaluation of properties. In addition, this approach can be used to inspect or sample the current state of the systems. This does not necessarily have to mean that the runtime verification algorithm is executed in the context of an application, but this event-generation can also be used to generate stimuli to external runtime verification engines that are implemented for example with the help of rewriting logic.

In the setting of field calculus programs, such an integration is more straight-forward: here, we do not need to establish a coupling between a target application and a runtime verification framework, but rather have FC programs that implement runtime verification monitors along side applications written in that formalism. As such, they use the same communication constructs to aggregate information from neighbours and trigger local actions.

As in the more traditional RV approaches for main-stream languages and systems, also here one can separate the implementation language from the specification language. We take a first step and show how common safety properties can be expressed as field calculus programs. Ideally, one would next strive for a specification language that resembles more a temporal logic with future or past operators, which is then translated into a field calculus program to monitor the property.

Taken as an approach to *distributed* runtime verification, we note that the field calculus also brings infrastructure that tackle a challenge in truly distributed systems: the dynamic nature of these systems with their varying number of participants and communication topology poses the challenge of reliability. So far, the RV community has mostly considered systems with a fixed number of agents and a fixed topology where communication is either point-to-point, allowing for interesting schemes to convey partial information, or broadcast, where message loss is not taken into account. See Basin *et al.*'s work [8] for a rare take on distributed runtime verification in the presence of communication delays and errors.

In field calculus, on the one hand one faces the same challenges, e.g. of establishing a global property across all agents. On the other hand, the constructs and mechanism of the field calculus, do provide a solution in themselves and do not require another level of middleware: a developer that is already familiar with the field calculus will naturally encode e.g. properties of resilience and awareness of network partitions into their specifications.

3 The field calculus

The *field calculus* [17] is a minimal 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 compose such distributed computations, by a level of abstraction that intentionally neglects explicitly management of synchronisation, message exchanges between devices, position and quantity of devices, and so on; while retaining Turing-universality for distributed computations [2].

3.1 The model of computation

In field calculus, a program P is periodically and asynchronously executed on every device, according to the following cyclic schedule. The involved device ι , every period T_ι :

1. perceives contextual information, which is formed by: data provided by sensors, local information stored in the previous round, and messages collected from neighbours while sleeping,¹ the latter in the form of a *neighbouring value* ϕ —essentially a map from neighbours to values v ;
2. evaluates the program P , considering as input the contextual information gathered as described above;
3. the result of this computation is a data structure that is stored locally, broadcast to neighbours, and possibly fed to actuators;
4. sleeps until it is awoken at the next activation.

By repetitive execution of such computation rounds, across space (where devices are located) and time (when devices fire), a global behaviour emerges [31], which can be fruitfully considered as occurring on the overall network of interconnected devices, modelled as a single aggregate machine equipped with a neighbouring-based topology relation. This process can be mathematically modelled through the notion of *event*, which correspond to the instants when devices are activated and start this sequence (see [2, 6] for further details on events and their role in modelling distributed computations).

Definition 1 (event [6]). An *event* e is modelled by the pair $e = (\iota, t)$ such that ι is the identifier of the device where the event takes place, and t is the time when the device ι is activated. The time stamp t refers to the local clock of ι .

Events are partially ordered by the following relationship.

Definition 2 (direct predecessor [6]). An event $e' = (\iota', t')$ is a *direct predecessor* (or *neighbour* for short) of an event $e = (\iota, t)$, denoted by $e' \rightsquigarrow e$, if the message broadcast by e' was the last from ι' able to reach ι before e occurred (and was not discarded by ι as an obsolete message).

It follows that if e' is a neighbour of e , then e' has to happen right before e , but not too long time ago (otherwise the message would have been discarded) or too far away (otherwise the message would not be received): thus, the neighbouring relation typically reflects spatial proximity. However, it could also be a logical relationship (e.g., connecting master devices to slave devices independently of their position), in which case the “far away” requirement would be measured through the logical network topology.

Furthermore, notice that the relation \rightsquigarrow on events forms a direct acyclic graph (DAG) among events, since cycles would correspond to a closed timelike curve. Hence, the \rightsquigarrow relation is time-driven and anti-symmetric, unlike spatial-only neighbouring (which is usually symmetrical).

¹Older messages may be retained until a certain timeout expires, or newer messages are received.

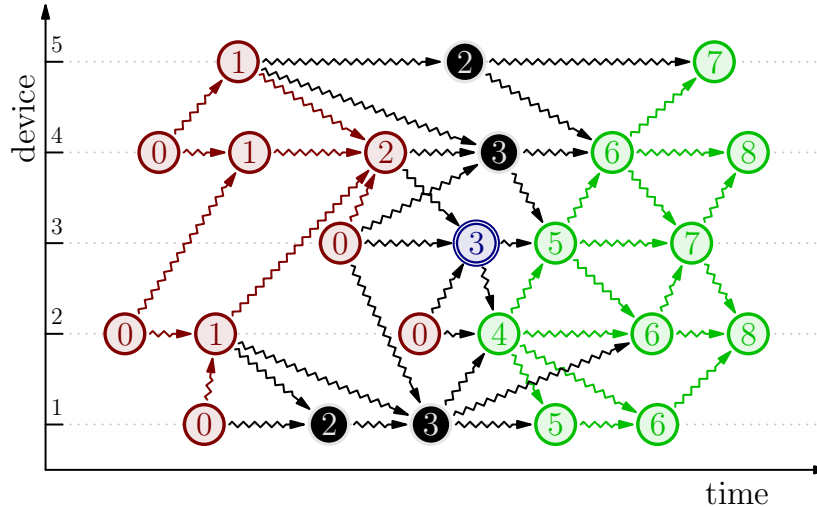


Figure 1: Representation of a field evolution of integers together with its underlying event structure (neighbouring). Past events of the circled blue event 3 are depicted in red, future events in green, concurrent events in black. This field evolution models the computation in each event of the longest preceding chain of events, obtainable locally by taking the maximum of the neighbour counters increased by 1.

Definition 3 (causality [6]). The *causality* partial order $e' \leq e$ on events is the transitive closure² of \rightsquigarrow .

The causality relation defines which events constitute the past, future or are concurrent to any given event. A set of events together with a neighbouring relation determines an *event structure*, represented in Figure 1. Notice that we do not assume that a global clock is available, nor that the scheduling of events follow a particularly regular pattern. This choice is dictated by the need to apply the field calculus to the broadest possible class of problems, as further restrictions can always be added without hassle.

Using the formalism of event structures, we can abstract the data manipulated by a field calculus program as a whole distributed space-time *field evolution* Φ , mapping individual events e in an event structure E to data values v (see Figure 1). Similarly, we can understand an “aggregate computing machine” as a device manipulating these field evolutions, and abstract is as a function mapping input field evolutions to output field evolutions.

3.2 The programming language

The syntax of field calculus is presented in Figure 2—the overbar notation \bar{e} is a shorthand for sequences of elements, and multiple overbars are intended to be expanded together, e.g., \bar{e} stands for e_1, \dots, e_n and $\bar{l} \mapsto \bar{\ell}$ for $l_1 \mapsto \ell_1, \dots, l_n \mapsto \ell_n$. The keywords `nbr` and `rep` correspond to the two peculiar constructs of field calculus, responsible of interaction and field dynamics, respectively; while `def` and `if` correspond to the standard function definition and the branching expression constructs.

A program P is the declaration of a set of functions \bar{F} of the kind “`def` $d(x_1, \dots, x_n) \{e\}$ ”, and a main expression e that is the one executed at each computation round, as well as the one considered (in the global viewpoint) as the overall field computation. An expressions e can be:

- A *variable* x , used e.g. as formal parameter of functions.
- A *value* v , which can be of the following two kinds:

²Thus, $e' \leq e$ iff there exists a (possibly empty) sequence e_1, \dots, e_n of events such that $e' \rightsquigarrow e_1 \rightsquigarrow \dots \rightsquigarrow e_n \rightsquigarrow e$.

P	$::= \bar{F} e$	program
F	$::= \text{def } d(\bar{x}) \{e\}$	function declaration
e	$::= x \mid f(\bar{e}) \mid v \mid \text{if}(e)\{e_1\}\{e_2\} \mid \text{nbr}\{e\} \mid \text{rep}(e)\{(x)=>e\}$	expression
f	$::= d \mid b$	function name
v	$::= \ell \mid \phi$	value
ℓ	$::= c(\bar{\ell})$	local value
ϕ	$::= \bar{t} \mapsto \bar{\ell}$	neighbouring field value

Figure 2: Syntax of the field calculus language.

- 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 devices t to local values ℓ , e.g., it could be the neighbouring value of distances of 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 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.
 - An *nbr-expression* $\text{nbr}\{e\}$, use to create a neighbouring field value mapping neighbours to their latest available result of evaluating e . In particular, each device t :
 1. broadcasts (together with its state information) its value of e to its neighbours,
 2. evaluates the expression into a neighbouring field value ϕ associating to each neighbour t' of t the latest evaluation of e at t' .

Note that the the evaluation by a device t of an *nbr-expression* within a branch of some $\text{if}(e_1)\dots$ expressions, is affected only by the neighbours of t that, during their last computation cycle, evaluated the same value for e_1 .

- A *rep-expression* $\text{rep}(e_1)\{(x)=>e_2\}$ models evolution through time, by returning the value of the expression e_2 where each occurrence of x is replaced by the value of the *rep-expression* at the previous computation cycle—or by e_1 if the *rep-expression* has not been evaluated in the previous computation cycle.

The meaning of a field calculus program can be defined through a denotational and an operational semantics, both thoroughly studied in [7]. The denotational semantics $\mathcal{E}[[e]]^E$ maps an expression e to a field evolution $\Phi = \mathcal{E}[[e]]^E$ on a given event structure E (see Section 3.1), and is compositional meaning that $\mathcal{E}[[f(e)]]^E = \mathcal{E}[[f]]^E(\mathcal{E}[[e]]^E)$.

Alternatively, an operational semantics of rounds in a network can be given in terms of a transition system $N \xrightarrow{act} N'$ between network configurations N , where *act* is either *env* to model any environment

change, or a device identifier to represent a device computation. These computations are in turn modelled by a local judgement $t; \Theta; \sigma \vdash e \Downarrow \theta$ to be read as “expression e evaluates to θ on device t with respect to sensor values σ and neighbours’ data Θ ”, where θ is the structure of values obtained from the evaluation of every sub-expression of e , and Θ is a map $\bar{t} \mapsto \bar{\theta}$ from neighbour device identifiers to the last θ_i which was received from them by the current device.

Example 1 (hop-count distance). In order to give an intuition of the behaviour of a field calculus program, consider the following function, where `minHood` selects the minimum element in the range of a numeric field ϕ , and `mux` is a classic “multiplexer” operator selecting its second or third argument depending on the truth value of the first (overloaded to apply pointwise on fields).

```
def hopcount(source) {
  rep (infinity) { (c) => mux(source, 0, minHood(nbr{c+1})) }
}
```

The `hopcount` function computes the number of hops required to reach a node where `source` is true: it is zero in sources, and equal to the minimum count of a neighbour incremented by one in non-source nodes.

Remark 2 (sample code). In practical implementations of the field calculus, the language is often extended to include additional features improving code readability. In Section 5 we shall use some of them, in particular:

- The traditional `let x = e_1 in e_2` construct, which can be thought as a shorthand for the expression `f(e_1, y_1, ..., y_n)` given the definition `def f(x, y_1, ..., y_n) {e_2}`, where y_1, \dots, y_n are the variables occurring free in e_2 .
- The notation `[e_1, ..., e_n]`, representing tuple creation `Tuple(e_1, ..., e_n)`.
- The multi-valued `rep` construct `rep (v_1, ..., v_n) {(x_1, ..., x_n) => e_1, ..., e_n}`, as a shorthand for the following.

```
rep ([v_1, ..., v_n]) { (t) =>
  let x_1 = 1st(t) in ... let x_n = nth(t) in [e_1, ..., e_n]
}
```

4 Implementing monitors in field calculus

Inspired by Francalanza et al. [20], we frame our discussion by considering a distributed monitoring setting where:

1. The system under analysis comprises a number of subsystems, identified by *processes* Π , that execute independently and might interact (i.e., synchronize or communicate) via the underlying communication platform.
2. The set of processes is partitioned across locations λ , i.e., each process Π is located at exactly one location, denoted by $loc(\Pi)$. Two processes Π and Π' are *local* to one other if and only if $loc(\Pi) = loc(\Pi')$, and *remote* otherwise. Processes may interact with both local and remote processes (usually remote communication is more expensive than local communication). Notable cases are when one of the following two conditions holds:
 - (a) There is just one location (i.e., all the processes are local);

- (b) At each location there is exactly one process (i.e., all processes are remote).
3. Each location hosts a number of *local traces* τ , each trace consists of a total ordered set of *events*, and each event describes a discrete computational step of a process at the location that hosts the trace. A trace may contain events of different process. Notable cases are when one or two of the following conditions holds:
 - (a) each trace contains events of a single process (i.e., each trace belongs to a single process), or
 - (b) for each process there is exactly one trace (containing the events of the process).
 4. Monitoring is performed by computation entities, identified by *monitors* M , that check properties of the system under analysis by analysing the traces. Similar to processes each monitor is hosted at a given location and may communicate with other (local or remote) monitors. Notable cases are when there is exactly one monitor for each:
 - (a) location,
 - (b) process, or
 - (c) trace.

In runtime verification monitors are generated from formal specifications. In the following we illustrate, by means of simple examples, how the field calculus can be used to implement distributed monitors. Our aim is to pave the way towards generating field calculus distributed monitors from suitable formal specifications. We consider the following setting:

- Each monitor is implemented by a field calculus program running on a dedicated (virtual or physical) device.
- Each local trace is mapped to a sensor.
- The `nbr` construct comes in two forms:
 - `nbrLocal`, for communication with local devices (i.e., if t_1 is a neighbour of t_2 then they are at the same location), and
 - `nbrRemote`, for communication with remote devices (i.e., if t_1 is a neighbour of t_2 then they are at different locations).
- Each device t is awoken whenever:
 - a new event arrives on one of the sensors of the devices, or
 - a new *different* message arrives from a (local or remote) neighbour t' ;³
 provided that a minimum time span T has elapsed from the previous evaluation cycle.

Moreover, for simplicity, we also assume that both conditions 3.a and 3.b hold. In the next two subsections we present examples in the context of the “local monitor only” and of the “remote monitors only” assumptions, respectively.

4.1 Local monitors only

In this subsection we assume that condition 2.a (given at the beginning of Section 4) holds, that is, every process is local. We consider two smart home scenarios, in which processes are assumed to be local through either: (i) physically wired connections; or (ii) short-range efficient wireless communication, as

³Note that if the new message is equal to the last message received from t' then the device t is not awoken.

the one expected by upcoming 5G standards. In this setting, the network topology can be full, that is, every node communicates with every other node.

In the first scenario, we want to monitor the following property: *air conditioning and lights are on when the room is not empty*. In order to express this property, we assume that the following 0-ary built-in operators (with corresponding traces) are given:

- `lights`: an optional Boolean value, which is true if the lights are on, false if they are off and null in nodes not controlling the lights.
- `people`: an optional Boolean value, depending on whether the node is sensing the presence of nearby people (if sensing is available).

This first property can then be expressed through the following program:

```
lights() == null || lights() == anyHood(nbrLocal{people() == true})
```

where `anyHood` is a built-in function that given a Boolean field ϕ , returns true if and only if at least one element in the range of ϕ is true. The monitoring property holds in nodes not controlling lights (i.e., when `lights` is `null`), or when the lights are on if and only if `people` is true in some communicating node, capturing the required idea.

In the second scenario, we want to monitor the following property: *if the volume of the stereo is above a certain threshold, every node should rapidly agree on alerting the stereo to lower its volume*. In order to express this property, we assume that the following 0-ary built-in operators (with corresponding traces) are given:

- `level`: the volume level of the stereo, or 0 in nodes not controlling the stereo.
- `alert`: an optional Boolean value, depending on whether the node is sensing excessive noise, which is null if no sensing is available.

This second property can then be expressed through the following program:

```
def roundsince(condition) {
  rep (0) { (x) => if (condition) {0} {x+1} }
}
roundsince(allHood(nbrLocal{alert() != false}) || level() <= THRESHOLD) < DELAY
```

where `THRESHOLD`, `DELAY` are given constants and `allHood` is a built-in function that given a Boolean field ϕ , returns true if and only if every element in the range of ϕ is true. Function `roundsince` counts the number of rounds elapsed since the last time `condition` was true. The monitoring property holds provided that no more than `DELAY` turns elapsed since when the volume was below `THRESHOLD` or all nodes agreed on alerting.

4.2 Remote monitors only

In this subsection we assume that condition 2.b (given at the beginning of Section 4) holds, that is, every process is remote. In this case, it is no longer realistic to assume a full communication topology; instead, we shall have few neighbours for every node to reduce the number of needed communications. This may not make a difference in case the property to monitor is fully local, as by the first example discussed in Section 2.1 which may be written through the following specification, where `nbrVector` is a returns the field of vectors to neighbours, `direction` is the quantity to be monitored and `angle` computes the relative angle between two vectors.

```
allHood(-60 < angle(nbrVector(), direction()) < 60 &&
        -60 < angle(-nbrVector(), nbr{direction()}) < 60)
```

However, when properties to monitor are not fully local, we may require a “data collection” routine to ensure effective spatial quantification (e.g., checking whether a property is true for all devices). This can be accomplished in field calculus through the *collection* building block, which is here instantiated for spatial quantification with the help of the result count of the simple distance estimation routine *hopcount* described in Example 1.

```
def everywhere(property, count) {
  rep (false) { (p) =>
    allHood(mux(nbrRemote{count} > count, nbrRemote{p}, property))
  } }
def somewhere(property, count) {
  rep (false) { (p) =>
    anyHood(mux(nbrRemote{count} > count, nbrRemote{p}, property))
  } }
```

The *everywhere* and *somewhere* functions check the validity of a property in nodes with a higher count, so that their value in the source should correspond to the intended result. More efficient collection [3, 31] and distance computation algorithms [4, 5] may be used in practical systems to implement those same functions: in this paper, we opted for the simplest implementations instead for sake of readability.

With the help of those functions, we can translate both scenarios in Section 4.1 to a remote-only setting. For the first scenario, we may want to check that an electronic system is on when some people are present in a large building, which can be accomplished by the following code.

```
lights() == null || lights() == somewhere(people() == true, hopcount(lights() != null))
```

For the second scenario, we may want to check that every area of such a building is alerted for evacuation after some dangerous event has been detected, which can be accomplished by the following code.

```
roundsince(everywhere(alert() != false, hopcount(level() != 0)) ||
           level() <= THRESHOLD) < DELAY
```

In both scenarios, we compute hop-count distances from controller nodes (which are reasonably unique), and use these distances to guide aggregation.

5 Monitoring field calculus programs

In case the distributed program to be monitored is a field calculus program, further opportunities arise from the ability of instrumenting the monitor code within the original algorithm, and possibly implementing feedback loops between them. Inspired by the second motivating example presented in Section 2.1, we consider the following *channel* routine building on the *hopcount* function presented in Example 1.

```
def broadcast(value, count) {
  rep (value) { (oldval) =>
    mux( count == 0, value, 2nd(minHood(nbr{[count, oldval]})) )
  } }
```

```

def elliptic-channel(sourcecount, destcount, width) {
  let sourcedest = broadcast(sourcecount, destcount) in
  sourcecount + destcount <= sourcedest + width
}
def channel(value, source, dest, width) {
  let sourcecount = hopcount(source) in
  let destcount = hopcount(dest) in
  let inarea = elliptic-channel(sourcecount, destcount, width) in
  if (inarea) { broadcast(value, sourcecount) } { value }
}

```

The broadcast function spreads a value from a source generating a certain hop-count distance (count) outwards: every device selects the provided value only if it is the source (count == 0), otherwise it selects the value of the neighbour with minimal count. Function elliptic-channel defines a roughly elliptic area with foci in a source and destination and given width, by comparing the sum of distances from the current location to the source and destination with the distance between the source and destination themselves (obtained by broadcasting from the destination the value of the distance to the source). Finally, function channel uses the above functions to broadcast a value in the area selected by elliptic-channel.

In order for the communication to be reliably performed, the width parameter has to be carefully tuned, depending also on the network characteristics. Thus, it is crucial to monitor the effectiveness of the choice, as performed by the following functions, where sumHood computes the sum of a numeric field ϕ , min computes the minimum between two numbers, and myID returns the identifier of the current device.

```

def samevalue(value, count) {
  let num,id = rep (1,myID()) { (num,id) =>
    sumHood(mux(nbr{id} == myID(), num, 0))+1,
    2nd(minHood( mux(nbr{value} == value, nbr{[count,myID()]}, [infinity, myID()]) ))
  } in
  broadcast(num, if (id == myID()) {0} {count})
}
def monitor(sourcecount, destcount, minw, maxw) {
  let w = min(samevalue(sourcecount,destcount), samevalue(destcount,sourcecount)) in
  if (w > maxw) {HIGH} {if (w < minw) {LOW} {OK}}
}

```

Function samevalue computes the number of devices holding the same value for value in devices with the lowest possible count: every device collects partial estimates num from neighbours who selected it in id, and selects in id the neighbour with the same value and lowest possible count. The num computed by the device with the lowest possible count is then broadcast to others (since devices with lowest possible count select themselves as id). The monitor then uses function samevalue to estimate the cross-section from both points of view of the source and destination, considering the minimum among them: a status is finally returned depending on whether this estimates fall above, below or within the required interval.

This monitor, if run within the area selected by elliptic-channel, can estimate whether the channel is properly established. Furthermore, it can be instrumented within the channel function to obtain an auto-adjusting channel as in the following.

```

def adjusting-channel(value, source, dest, minw, maxw) {
  let sourcecount = hopcount(source) in
  let destcount   = hopcount(dest) in
  let inarea = 1st(rep (False, maxw) { (oarea, owidth) =>
    let narea = elliptic-channel(sourcecount, destcount, owidth) in
    let status = if (narea) {monitor(narea, minw, maxw)} {OK} in
    narea, if (status == OK) {width} {if (status == LOW) {owidth+1} {owidth-1}}
  }) in
  if (inarea) { broadcast(value, sourcecount) } { value }
}

```

This function increases or decreases the width by 1 according to the status returned by the monitor. Furthermore, it does so independently in *every device* of the network, allowing the shape of the channel to adjust to the network local peculiarities (instead of the fixed elliptical shape of the traditional channel).

6 Conclusion

In this position paper we have illustrated, by means of simple examples, how the field calculus can be used to implement distributed monitors in different settings. In particular, we have provided examples of local and remote monitors, and an example of a field calculus program within which the monitor can be instrumented providing the algorithm with an additional auto-correcting power.

In future work we would like to investigate how field calculus expressions, e.g. using the `nbr-construct`, could be used in conjunction with a specification language like LTL; and possibly be automatically generated by a logical language. This would allow us to write properties along the lines of “Eventually, all my neighbours...” or “Some neighbour will always...”.

Acknowledgements

This work has been partially supported by the European Union’s Horizon 2020 research and innovation programme under project COEMS (www.coems.eu, grant agreement no. 732016), project Hy-Var (www.hyvar-project.eu, grant agreement no. 644298) and ICT COST Action IC1402 ARVI (www.cost-arvi.eu). We thank the anonymous VORTEX 2018 reviewers for insightful comments and suggestions for improving the presentation.

References

- [1] Luca Aceto, Antonis Achilleos, Adrian Francalanza & Anna Ingólfssdóttir (2018): *A Framework for Parameterized Monitorability*. In: *Foundations of Software Science and Computation Structures, Lecture Notes in Computer Science* 10803, Springer, pp. 203–220, doi:10.1007/978-3-319-89366-2_11.
- [2] Giorgio Audrito, Jacob Beal, Ferruccio Damiani & Mirko Viroli (2018): *Space-Time Universality of Field Calculus*. In: *Coordination Models and Languages, Lecture Notes in Computer Science* 10852, Springer, pp. 1–20, doi:10.1007/978-3-319-92408-3_1.
- [3] Giorgio Audrito & Sergio Bergamini (2017): *Resilient Blocks for Summarising Distributed Data*. In: *Proceedings of ALP4IoT@iFM 2017, EPTCS* 264, pp. 23–26, doi:10.4204/EPTCS.264.3.
- [4] Giorgio Audrito, Roberto Casadei, Ferruccio Damiani & Mirko Viroli (2017): *Compositional Blocks for Optimal Self-Healing Gradients*. In: *11th IEEE International Conference on Self-Adaptive and Self-Organizing Systems*, IEEE Computer Society, pp. 91–100, doi:10.1109/SASO.2017.18.

- [5] Giorgio Audrito, Ferruccio Damiani & Mirko Viroli (2018): *Optimal single-path information propagation in gradient-based algorithms*. *Sci. Comput. Program.* 166, pp. 146–166, doi:10.1016/j.scico.2018.06.002.
- [6] Giorgio Audrito, Ferruccio Damiani, Mirko Viroli & Enrico Bini (2018): *Distributed Real-Time Shortest-Paths Computations with the Field Calculus*. In: *2018 IEEE Real-Time Systems Symposium (RTSS)*, pp. 23–34, doi:10.1109/RTSS.2018.00013.
- [7] Giorgio Audrito, Mirko Viroli, Ferruccio Damiani, Danilo Pianini & Jacob Beal (2019): *A Higher-Order Calculus of Computational Fields*. *ACM Trans. Comput. Logic* 20(1), pp. 5:1–5:55, doi:10.1145/3285956.
- [8] David A. Basin, Felix Klaedtke & Eugen Zalinescu (2015): *Failure-aware Runtime Verification of Distributed Systems*. In Harsha & Ramalingam [22], pp. 590–603, doi:10.4230/LIPIcs.FSTTCS.2015.590.
- [9] Andreas Bauer & Yliès Falcone (2016): *Decentralised LTL monitoring*. *Formal Methods in System Design* 48(1-2), pp. 46–93, doi:10.1007/s10703-016-0253-8.
- [10] Jacob Beal, Stefan Dulman, Kyle Usbeck, Mirko Viroli & Nikolaus Correll (2013): *Organizing the Aggregate: Languages for Spatial Computing*. In: *Formal and Practical Aspects of Domain-Specific Languages: Recent Developments*, chapter 16, IGI Global, pp. 436–501, doi:10.4018/978-1-4666-2092-6.ch016. A longer version available at: <http://arxiv.org/abs/1202.5509>.
- [11] Jacob Beal, Danilo Pianini & Mirko Viroli (2015): *Aggregate Programming for the Internet of Things*. *IEEE Computer* 48(9), pp. 22–30, doi:10.1109/MC.2015.261.
- [12] Jacob Beal, Mirko Viroli, Danilo Pianini & Ferruccio Damiani (2017): *Self-Adaptation to Device Distribution in the Internet of Things*. *ACM Transaction on Autonomous and Adaptive Systems* 12(3), pp. 12:1–12:29, doi:10.1145/3105758.
- [13] Nicola Biccocchi, Marco Mamei & Franco Zambonelli (2012): *Self-organizing virtual macro sensors*. *TAAS* 7(1), pp. 2:1–2:28, doi:10.1145/2168260.2168262.
- [14] Roberto Casadei & Mirko Viroli (2016): *Towards Aggregate Programming in Scala*. In: *First Workshop on Programming Models and Languages for Distributed Computing, PMLDC '16*, ACM, New York, NY, USA, pp. 5:1–5:7, doi:10.1145/2957319.2.
- [15] Daniel Coore (1999): *Botanical Computing: A Developmental Approach to Generating Inter connect Topologies on an Amorphous Computer*. Ph.D. thesis, MIT, Cambridge, MA, USA. Available at <http://hdl.handle.net/1721.1/80483>.
- [16] Carlo Curino, Matteo Giani, Marco Giorgetta, Alessandro Giusti, Amy L. Murphy & Gian Pietro Picco (2005): *Mobile data collection in sensor networks: The TinyLime*. *Pervasive and Mobile Computing* 1(4), pp. 446–469, doi:10.1016/j.pmcj.2005.08.003.
- [17] Ferruccio Damiani, Mirko Viroli & Jacob Beal (2016): *A type-sound calculus of computational fields*. *Science of Computer Programming* 117, pp. 17–44, doi:10.1016/j.scico.2015.11.005.
- [18] Jose Luis Fernandez-Marquez, Giovanna Di Marzo Serugendo, Sara Montagna, Mirko Viroli & Josep Lluís Arcos (2013): *Description and composition of bio-inspired design patterns: a complete overview*. *Natural Computing* 12(1), pp. 43–67, doi:10.1007/s11047-012-9324-y.
- [19] Bernd Finkbeiner, Christopher Hahn, Marvin Stenger & Leander Tentrup (2017): *Monitoring Hyperproperties*. In Lahiri & Reger [23], pp. 190–207, doi:10.1007/978-3-319-67531-2_12.
- [20] Adrian Francalanza, Jorge A. Pérez & César Sánchez (2018): *Runtime Verification for Decentralised and Distributed Systems*. In Ezio Bartocci & Yliès Falcone, editors: *Lectures on Runtime Verification: Introductory and Advanced Topics, Lecture Notes in Computer Science* 10457, Springer, pp. 176–210, doi:10.1007/978-3-319-75632-5_6.
- [21] Jean-Louis Giavitto, Olivier Michel, Julien Cohen & Antoine Spicher (2004): *Computations in Space and Space in Computations*. In: *Unconventional Programming Paradigms, Lecture Notes in Computer Science* 3566, Springer, pp. 137–152, doi:10.1007/11527800_11.

- [22] Prahlahd Harsha & G. Ramalingam, editors (2015): *35th IARCS Annual Conference on Foundation of Software Technology and Theoretical Computer Science, FSTTCS 2015. LIPIcs 45*, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, doi:10.4230/LIPIcs.FSTTCS.2015.i.
- [23] Shuvendu K. Lahiri & Giles Regeer, editors (2017): *Runtime Verification - 17th International Conference, RV 2017, Seattle, WA, USA, September 13-16, 2017, Proceedings. Lecture Notes in Computer Science 10548*, Springer, doi:10.1007/978-3-319-67531-2.
- [24] Martin Leucker & Christian Schallhart (2009): *A brief account of runtime verification. J. Log. Algebr. Program.* 78(5), pp. 293–303, doi:10.1016/j.jlap.2008.08.004.
- [25] Alberto Lluch-Lafuente, Michele Loreti & Ugo Montanari (2017): *Asynchronous Distributed Execution Of Fixpoint-Based Computational Fields. Logical Methods in Computer Science* 13(1), doi:10.23638/LMCS-13(1:13)2017.
- [26] Marco Mamei & Franco Zambonelli (2009): *Programming pervasive and mobile computing applications: The TOTA approach. ACM Trans. on Software Engineering Methodologies* 18(4), pp. 1–56, doi:10.1145/1538942.1538945.
- [27] Ryan Newton & Matt Welsh (2004): *Region streams: functional macroprogramming for sensor networks. In: Workshop on Data Management for Sensor Networks, ACM International Conference Proceeding Series 72*, ACM, pp. 78–87, doi:10.1145/1052199.1052213.
- [28] Yuichi Nishiwaki (2016): *Digamma-Calculus: A Universal Programming Language of Self-Stabilizing Computational Fields. In: eCAS, Self-Adaptive and Self-Organizing Systems Workshops, IEEE*, pp. 198–203, doi:10.1109/FAS-W.2016.51.
- [29] K. Sen, A. Vardhan, G. Agha & G. Rosu (2004): *Efficient decentralized monitoring of safety in distributed systems. In: 26th Intl. Conf. on Software Engineering*, pp. 418–427, doi:10.1109/ICSE.2004.1317464.
- [30] Volker Stolz & Eric Bodden (2006): *Temporal Assertions using AspectJ. Electr. Notes Theor. Comput. Sci.* 144(4), pp. 109–124, doi:10.1016/j.entcs.2006.02.007.
- [31] Mirko Viroli, Giorgio Audrito, Jacob Beal, Ferruccio Damiani & Danilo Pianini (2018): *Engineering Resilient Collective Adaptive Systems by Self-Stabilisation. ACM Transactions on Modelling and Computer Simulation* 28(2), pp. 16:1–16:28, doi:10.1145/3177774.
- [32] Mirko Viroli, Jacob Beal, Ferruccio Damiani, Giorgio Audrito, Roberto Casadei & Danilo Pianini (2018): *From Field-Based Coordination to Aggregate Computing. In: Coordination Models and Languages, Lecture Notes in Computer Science 10852*, Springer, pp. 252–279, doi:10.1007/978-3-319-92408-3_12.
- [33] Kamin Whitehouse, Cory Sharp, David E. Culler & Eric A. Brewer (2004): *Hood: A Neighborhood Abstraction for Sensor Networks. In: 2nd International Conference on Mobile Systems, Applications, and Services, ACM / USENIX*, pp. 99–110, doi:10.1145/990064.990079.