

PROGRAMMING DISTRIBUTED COLLECTIVE PROCESSES IN THE EXCHANGE CALCULUS

GIORGIO AUDRITO ^a, ROBERTO CASADEI ^b, FERRUCCIO DAMIANI ^a,
GIANLUCA TORTA ^a, AND MIRKO VIROLI ^b

^a Università di Torino, Turin, Italy

e-mail address: giorgio.audrito@unito.it, ferruccio.damiani@unito.it, gianluca.torta@unito.it

^b Alma Mater Studiorum—Università di Bologna, Cesena, Italy

e-mail address: robby.casadei@unibo.it, mirko.viroli@unibo.it

ABSTRACT. Recent trends like the Internet of Things (IoT) suggest a vision of dense and multi-scale deployments of computing devices in nearly all kinds of environments. A prominent engineering challenge revolves around programming the collective adaptive behaviour of such computational ecosystems. This requires abstractions able to capture concepts like ensembles (dynamic groups of cooperating devices) and collective tasks (joint activities carried out by ensembles). In this work, we consider collections of devices interacting with neighbours and that execute in nearly-synchronised sense–compute–interact rounds, where the computation is given by a single program mapping sensing values and incoming messages to output and outgoing messages. To support programming whole computational collectives, we propose the abstraction of a distributed collective process, which can be used to define at once the ensemble formation logic and its collective task. We formalise the abstraction in the eXchange Calculus (XC), a core functional language based on neighbouring values (maps from neighbours to values) where state and interaction is handled through a single primitive, **exchange**, and provide a corresponding implementation in the FCPP language. Then, we exercise distributed collective processes using two case studies: multi-hop message propagation and distributed monitoring of spatial properties. Finally, we discuss the features of the abstraction and its suitability for different kinds of distributed computing applications.

1. INTRODUCTION

Programming the collective behaviour of large groups of interactive computing devices is a major research problem, urged by scenarios like the Internet of Things [AIM10] and swarm robotics [BFBD13]. This problem is (partially) addressed by several related research threads including *coordination* [GC92, PA98], *multi-agent systems* [BBHR20], *collective adaptive systems* [DJW20], *macroprogramming* [Cas23b, NW04], and *spatial computing* [BDU⁺13].

This activity can be supported by suitable *programming abstractions* supporting declarative specifications of collective behaviours. Examples of abstractions include

Key words and phrases: collective computing, collective processes, ensembles, formation control, multi-agent systems, programming languages and self-organisation.

ensembles [NLPT14], computational fields [VBD⁺19], collective communication interfaces [WM04, AADNL20], and collective-based tasks [SSV⁺20]. In this work, we cover the abstraction of a *distributed collective process (DCP)*, inspired by aggregate processes [CVA⁺19, CVA⁺21, TADT22], which provides a way to specify collective activities that run concurrently and spread out on evolving domains of devices.

In this paper, we interpret the DCP mechanism in the abstract model of event structures [NPW81, ABDV18], and present its realisation in terms of a programming construct in the eXchange Calculus (XC) [ACD⁺22]. XC is a core functional language that generalises field calculi [VBD⁺19] and promotes programming the self-organisation in systems of neighbour-interacting devices. The XC-based formalisation of DCPs is also implemented in the FCPP language [Aud20]. Then, we discuss the abstraction, and present examples of use of DCPs, showing they can capture several patterns of collective adaptive behaviour and self-organisation. We use the FCPP simulator [ART22] to assess that DCP-based XC programs provide the intended functionality.

In summary, we build on previous work on the field calculus and aggregate processes [CVA⁺21] (see Section 7 for details), and provide the following novel contributions:

- we propose the DCP abstraction as a generalisation of the field calculus-based mechanism of aggregate processes where process propagation can be controlled by the sender (and not just through the *opt-out* of the recipient);
- we formalise and implement the DCP abstraction in the XC formal language and the FCPP programming language;
- we use simulated case studies (publicly available at a GitHub repository—cf. Footnote 1) and examples to discuss the features of the DCP abstraction and show performance benefits w.r.t. aggregate processes.

This manuscript is an extended version of paper [ACD⁺23] presented at COORDINATION’23. Beyond such content, it includes a self-contained formalisation (Sections 3 and 4), a description of the implementation in the FCPP language (Section 4.4), a whole new set of simulation-based experiments (Sections 5.1 and 5.2), additional examples and discussion (Section 6), and a larger coverage of related work (Section 7).

The remainder of the paper is structured as follows. Section 2 positions the work, explicating its context and motivation. Section 3 provides background on XC. Section 4 provides the contribution: formalisation, description of the abstraction, and implementation. Section 5 presents two simulated use cases that exploit the proposed techniques, where quantitative metrics are extracted to assess correctness and resource usage. Section 6 provides a broader discussion of the proposed programming framework, together with further examples to highlight its features and its applicability. Section 7 covers related work. Section 8 provides a wrap-up with further insights on limitations to be addressed in future work.

2. CONTEXT AND MOTIVATION

This work lies in the context of models and languages for programming collective behaviour [DJW20, Cas23a, BFBD13]. Indeed, achieving the desired collective behaviour is an engineering goal for different domains and applications:

- *Swarm robotics*. Multiple robots may be tasked to move and act as a collective to explore an unknown environment [MWT⁺19], to search and rescue victims for humanitarian aid

- after disasters [AJAM20], to map a crop field for the presence of weeds [AIHT17], to transport objects whose weight exceeds the limit of individual robots [GD09], etc.
- *The IoT*. The *things* should coordinate to promote application-level goals (e.g., by gathering and processing relevant data) while making efficient use of resources. For instance, the nodes could support the aggregation of machine learning models [SYN⁺21], or collaborate to measure the collective status of the network to support various activities ranging from environment sensing [LHJ19] to remote attestation of system integrity [ACL⁺20].
 - *Hybrid Collective Intelligence (CI)*. Socio-technical systems involving humans and computing devices could be programmed as “social machines” [HB10] executing orchestrated tasks [SSV⁺20], or fostering the emergence of collective knowledge [Gru08].
 - *Computing ecosystems*. Modern infrastructures spanning the edge-cloud continuum can be conceived as collective systems. The computing nodes could share and process information to create efficient topologies [KS21, PCVN21] to coordinate task allocation and execution [MK16], resiliently.

This problem is at the core of several related research threads. The field of *coordination* [GC92, PA98] addresses it by governing interaction; *collective adaptive systems* engineering [DJW20] investigates means for collective adaptation in large populations of agents; *spatial computing* [BDU⁺13] leverages spatial abstractions to drive behaviour and perform computation about and *in* space; *macroprogramming* [Cas23b, NW04] takes a programming language-perspective to expressing and steering macroscopic behaviour; *multi-agent systems* [BBHR20] consider how autonomous entities can solve problems through cooperation, competition, cognitive mechanisms, organisational structures, etc.

In this work, we consider a language-based software engineering perspective [Gup15]. In the literature, several abstractions have been proposed: ensembles, collective communication interfaces, spatiotemporal regions, etc. (see Section 7 for details). Similarly, we seek for abstractions for expressing collective behaviour, with some requirements and reference contexts that we delineate in the following.

In particular, we consider *collective systems*, namely largely homogeneous collections of devices or agents. Each device can be thought of as a resource that provides capabilities and provides access to a local context that depends on its situation on the environment and possibly its state. For instance, in a smart city, fixed smart lights may be located nearby streets, smart cameras may support monitoring of facilities, smart vehicles may move around to gather city-wide infrastructural data, etc. To reduce bottlenecks and single-points-of-failure, we avoid centralisations and opt for a fully decentralised approach: a device may interact only within its local context (e.g., a portion of the environment and other nearby devices). If our goal is to exploit the distributed, pervasive computer made of an entire collection of situated devices, an idea could be to run *collaborative* tasks involving subsets of devices—to exploit their resources, capabilities, and contexts. Since a process may not know beforehand the resources/capabilities it needs and the relevant contexts, it may embed the logic to look for them, i.e., to spread over the collective system until its proper set of supporting devices have been identified. Moreover, the requirements of the process may change over time, dynamically self-adapting to various environmental conditions and changing goals. Within a process that concurrently spans a collection of devices, local computations may be scheduled and information may flow around in order to support collective activities [BFBD13, WG09] such as collective perception [GRWF16], collective decision-making [Bul14], collective movement [NM13], etc. So, if the collective that sustains the process decides that more resources are needed, the process may spread to a larger set

of devices; conversely, if the collective task has been completed, the devices may quit the process, eventually making it vanish. Informally, this is our idea of a *distributed collective processes (DCP)*: a *process* (i.e., a running program or task) which is *collective* (i.e., a joint activity carried out by groups of devices) and *distributed* (i.e., concurrently active on multiple devices), where the collective task and the underlying ensemble can mutually affect each other, and ensemble formation is driven by decentralised device-to-device interaction.

Approach Overview. In this paper, we aim to characterise the aforementioned abstract idea of a DCP in terms of the following main ingredients:

- *event structures*: this formal model will serve to represent collective system executions (i.e., networks of device computation rounds causally-related by message-passing) and to describe DCPs in a general setting;
- *eXchange Calculus (XC)*: this core language captures a minimal set of constructs for expressing collective computations (i.e., the activity to be performed at the different events of a collective system execution); we will extend the XC language to provide a *programming interface for DCPs*.

The following section provides background on these two ingredients.

3. BACKGROUND: THE EXCHANGE CALCULUS

We consider the *eXchange Calculus (XC)* [ACD⁺22] as the formal framework for modelling, reasoning about, and implementing DCPs. In this section, we start by providing an overview of the calculus (Section 3.1), then present the system and execution model (Section 3.2), providing an operational view of the kinds of systems we target, and later describe the basic constructs of XC that we leverage in this work (Section 3.3). Finally, we conclude with the formal details of XC: syntax (Section 3.4), typing (Section 3.5), and semantics (Section 3.6).

3.1. Overview. XC is a core functional language capturing a minimal set of constructs for programming the self-organising behaviour of collections of neighbour-interacting devices (as in robot swarms, sensor networks, and IoT systems). It assumes a system model where a collection of devices (equipped with sensors and actuators for interacting with their environment) operate in asynchronous *sense–compute–interact* rounds. An XC program expresses the logic at the “compute” step: it receives as input the local context of the device (i.e., sensor values and data from neighbours) and returns as output the actuations to perform and the data to be shared with neighbours.

Structurally, XC is basically a lambda calculus extended with the *neighbouring value (nvalue)* data type and a single communication construct called *exchange*. XC generalises over field calculi [VBD⁺19, ACDV23], which historically had (i) no support for sending different messages to different neighbours, and (ii) different constructs for state management and communication. An nvalue is a map \underline{w} from device identifiers δ_i to local values ℓ_i , with an additional *default* value ℓ for missing keys, written as $\underline{w} = \ell[\delta_1 \mapsto \ell_1, \dots, \delta_n \mapsto \ell_n]$. Construct **exchange**($\mathbf{e}_i, (\underline{n}) \Rightarrow (\mathbf{e}_r, \mathbf{e}_s)$) evolves a distributed quantity in space and time by mapping the nvalue of neighbour messages \underline{n} (using the value from \mathbf{e}_i as default) to a return value (computed in expression \mathbf{e}_r) and a value to be sent (expression \mathbf{e}_s).

To formalise XC, it is typical (as for other field calculi) to distinguish the overall system execution (i.e., environment evolution, exchange of messages, scheduling of rounds) from the execution of individual rounds. Accordingly, XC is formalised in terms of (i) an operational

semantics describing how a device interprets an XC program, and (ii) a network semantics expressing the relationships between rounds, in the framework of *event structures*.

XC can succinctly encode basic self-organisation building blocks, which can in turn be composed to implement more complex collective behaviours. For instance, a gradient, namely the program computing distance estimates from each device to their closest source device, can be re-used to provide optimal directions for propagating and collecting distributed information. The ability to program emergent behaviour based on repeated execution of programs supporting the processing and diffusion of local information is especially instrumental to build DCPs—as it will be covered in Section 4.

3.2. System Model. The target system that we would like to program can be modelled as a collection of *nodes*, able to interact with the environment through *sensors* and *actuators*, and able to communicate with *neighbours* by exchanging messages. We assume that each node runs in asynchronous *sense–compute–act rounds*, where

- (1) *sense*: the node queries sensors for getting up-to-date environmental values, and gathers recent messages from neighbours (which may expire after some time)—all this information is what we call as the node’s *context*;
- (2) *compute*: the node evaluates the common control program, mapping the context (i.e., inputs from sensors and neighbours) to an output describing the actions to perform (i.e., actuations and communications);
- (3) *act*: the node executes the actions as dictated by the program, possibly resulting into environment change or message delivery to neighbours.

This kind of loop is used to ensure that the context is continuously assessed (at discrete times), and the reactions are also computed and performed continuously. This model has shown to be particularly useful to specify self-organising and collective adaptive behaviours, especially for long-running coordination tasks [VBD⁺19]. Note that details regarding aspects like the topology (as defined by a neighbouring relationship), round scheduling, message expiration, and how actual communication is carried out can be defined by middleware implementations and also tuned on a per-application basis. What is key here is the “repeated” evaluation and action on the local context by each device; upon such change acquired by devices, computations regulate how the novel information is integrated with past information, and how all this is transformed and shared with neighbours. Flowing from neighbourhood to neighbourhood, the effects of computation would *eventually spread and stabilise* [VAB⁺18] *on a global scale* (cf. *emergence* [WH07]). So, in this work, our focus is on computations that are long-running, “progressive”, and collective (based on collaboration of groups of devices).

The semantics of a system execution can be expressed as an event structure (see Figure 1), where events ϵ denote whole sense–compute–act rounds, and arrows between events denote that certain source events have provided inputs (i.e., messages) to target events. In particular, if event ϵ' is connected with an arrow to ϵ , we say that ϵ' is a *neighbour* of ϵ , denoted $\epsilon' \rightsquigarrow \epsilon$. We denote with $N(\epsilon)$ the set of all *neighbours* of ϵ , with $d(\epsilon)$ the device where event ϵ happens (i.e., where it is executed), and with $p(\epsilon)$ the empty or singleton set with only the previous event of ϵ at the same device.

Programming such systems ultimately means defining the control rules that specify how the local context at each event is mapped to the messages to be sent to neighbour events.

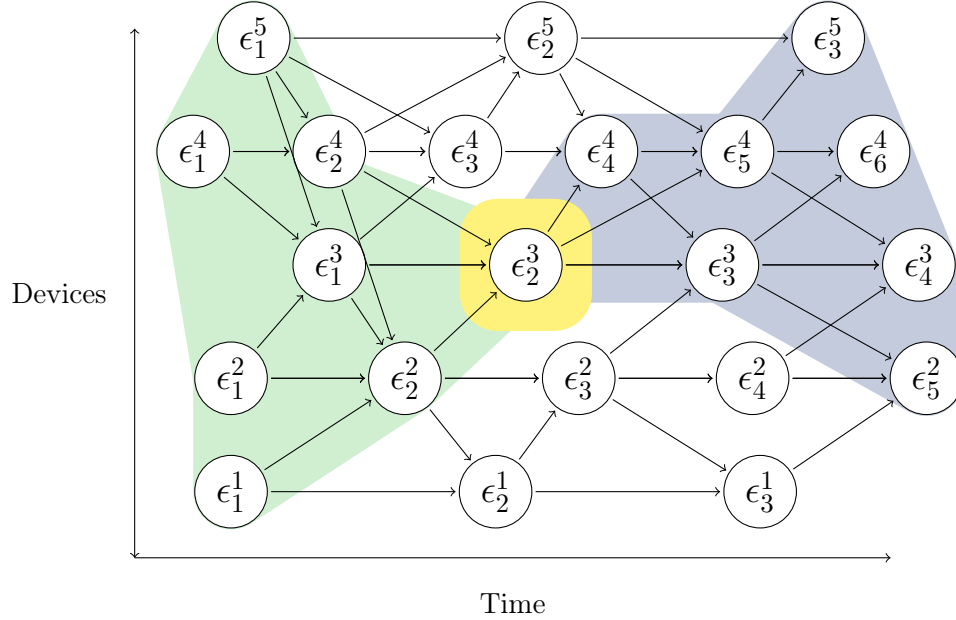


Figure 1: Example of an event structure modelling a distributed system execution. Nodes labelled by ϵ_k^δ denote the k -th round of device δ . The yellow background highlights a reference event, from which its past (green) and future (blue) are identified through the causal relationship implied by the arrows denoting neighbour events.

3.3. XC key data type: Neighbouring Values. In XC, we distinguish two types of values. The *Local* values ℓ include classic atomic and structured types A such as int, float, string, and list. The neighbouring values (*nvalues*) are instead maps \underline{w} from device identifiers δ_i to corresponding local values ℓ_i , with an additional local value ℓ that acts as a *default*:

$$\underline{w} = \ell[\delta_1 \mapsto \ell_1, \dots, \delta_n \mapsto \ell_n]$$

A *nvalue* specifies a (set of) values received from or sent to neighbours: received values are gathered into *nvalues*, then can be locally processed, and the final resulting *nvalue* can be interpreted as messages to be sent back to neighbours. The devices with an associated entry in the *nvalue* are thus typically a (small) subset of all devices, namely those that are close-enough to the current device, and which are assumed to work correctly.

The default is used when a value is not available for some neighbour δ' , e.g., because δ' has just been switched on and has not yet produced a value, or because it has just moved close enough to the current device δ to become one of its neighbours. The notation above should therefore read as “the *nvalue* \underline{w} is ℓ everywhere (i.e., for all neighbours) except for devices $\delta_1, \dots, \delta_n$ which have values ℓ_1, \dots, ℓ_n ”.

To exemplify *nvalues*, in Figure 1, upon waking up for computation ϵ_3^3 , device δ_3 may process a *nvalue* $\underline{w} = 0[\delta_4 \mapsto 1, \delta_3 \mapsto 2, \delta_2 \mapsto 3]$, corresponding to the messages carrying scalar values 1, 2, and 3 received when asleep from δ_4, δ_3 (i.e., *itself* at the previous round), and δ_2 . The entries for all other (neighbour) devices default to 0. After the computation, δ_2 may send out the messages denoted by the *nvalue* $\underline{w}' = 0[\delta_4 \mapsto 5, \delta_3 \mapsto 6]$; so that 5 is sent to δ_4 , 6 is sent to δ_3 , and 0 is sent to every other device, such as a newly-connected device. For

convenience, we may use the notation $\underline{w}(\delta')$ for the local value (specific or default) associated with δ' by \underline{w} .

Note that a local value ℓ can be naturally converted to a nvalue $\ell[]$ where it is the default value for every device. Except for clarity, thus local values and nvalues can be treated uniformly. Functions on local values are implicitly lifted to nvalues, by applying them on the maps' content pointwise. For example, if \underline{w}_1 assigns value 2 to δ_3 and \underline{w}_2 assigns default value 1 to δ_3 , then $\underline{w}_3 = \underline{w}_1 \cdot \underline{w}_2$ shall assign value $2 \cdot 1 = 2$ to δ_3 . In XC, we also lift to nvalues the *multiplexer* operator $\text{mux}(\ell_1, \ell_2, \ell_3)$, which returns ℓ_2 if ℓ_1 is `True` and ℓ_3 otherwise.

A fundamental operation on nvalues is provided by the built-in function $\text{ifold}(f : (A, B) \rightarrow A, \underline{w} : \underline{B}, \ell : A) : A$. As suggested by the name, the function folds over a nvalue \underline{w} , i.e., starting from a base local value ℓ it repeatedly applies function f to neighbours' values in \underline{w} , excluding the value for the current device. For instance, if δ_2 with set of neighbours $\{\delta_1, \delta_3, \delta_4\}$ performs a `ifold` operation $\text{ifold}(*, \underline{w}, 1)$, the output will be $1 \cdot \underline{w}(\delta_1) \cdot \underline{w}(\delta_3) \cdot \underline{w}(\delta_4)$. Note that, as nvalues are *unordered maps*, it is sensible to assume that f is associative and commutative.

Two built-ins on nvalues act on the value associated with the current (self) device:

- `self`($\underline{w} : \underline{A}$) : A returns the local value $\underline{w}(\delta)$ in \underline{w} for the self device δ
- `updateSelf`($\underline{w} : \underline{A}, \ell : A$) : \underline{A} returns an nvalue where the value for δ is set to ℓ .

Furthermore, XC features a single communication primitive `exchange`($\mathbf{e}_i, (\underline{n}) \Rightarrow (\mathbf{e}_r, \mathbf{e}_s)$) which is evaluated as follows.

- the device computes the local value ℓ_i of \mathbf{e}_i (the *initial* value);
- to evaluate the function provided as second argument, it substitutes variable \underline{n} with the nvalue \underline{w} of messages received from the neighbours for this exchange, using ℓ_i as default;
- the exchange expression as a whole returns the (neighbouring or local) value v_r from the evaluation of \mathbf{e}_r ;
- the second expression of the argument function's output, \mathbf{e}_s , evaluates to a nvalue \underline{w}_s consisting of local values to be sent to neighbour devices δ' , that will use their corresponding $\underline{w}_s(\delta')$ as soon as they wake up and perform their next execution round.

Since expressions v_r and v_s often coincide, we also provide `exchange`($\mathbf{e}_i, (\underline{n}) \Rightarrow \text{retsend } \mathbf{e}$) as a shorthand for `exchange`($\mathbf{e}_i, (\underline{n}) \Rightarrow (\mathbf{e}, \mathbf{e})$). This construct abstracts the general concept of message exchange, and is sufficiently expressive to allow other communication patterns to be expressed by it. If a program executes multiple exchange-expressions, XC ensures through *alignment* that the messages are dispatched to corresponding exchange-expressions across rounds. The detailed workings of this process will be given in Section 3.6.

Example 3.1 (Gradient). A simple computation that can be executed on event structures is a distributed gradient algorithm, which computes everywhere in a network of devices the minimum distance from each device to its closest *source* device.

The *Adaptive Bellman-Ford (ABF)* algorithm [MDB19] computes the gradient by assuming that each device is able to get (an estimation of) its distance to each one of its neighbours. Let S denote the set of sources, $D_\epsilon(\delta')$ be the distance between devices $\delta = d(\epsilon)$ and δ' as estimated in event ϵ , and $g(\epsilon)$ be the resulting gradient value computed at ϵ . Then, the ABF algorithm on event structures can be expressed as follows:

$$g(\epsilon) = \begin{cases} \min_{\epsilon' \in N(\epsilon) \setminus p(\epsilon)} \{g(\epsilon') + D_\epsilon(d(\epsilon'))\} & d(\epsilon) \notin S \\ 0 & d(\epsilon) \in S \end{cases}$$

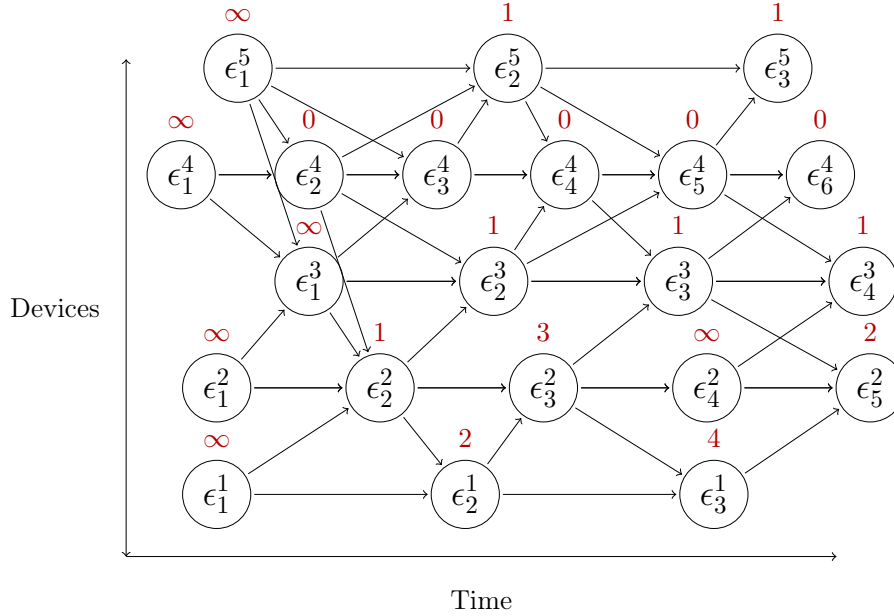


Figure 2: Example of a gradient computation (cf. Example 3.1) on the same event structure of Figure 1. Assume that device 4 becomes source at its 2nd round, i.e., at event ϵ_2^4 , and that neighbouring events are at a unit distance, i.e., $D_\epsilon(d(\epsilon')) = 1$ for all $\epsilon \rightsquigarrow \epsilon'$. Then, the output of the gradient computation at some event is shown with a red label above that event. In each event, the gradient value is computed by taking the gradient values at past neighbour events (excluding past events of the same device), adding the corresponding distance (which we assume any device is able to estimate locally), and selecting the minimum of them. For instance, ϵ_5^2 gathers [$\epsilon_3^3 \mapsto 1, \epsilon_3^1 \mapsto 4$], adds neighbouring distances to get [$\epsilon_3^3 \mapsto 2, \epsilon_3^1 \mapsto 5$], and selects the minimum value, i.e., 2. If there are no neighbouring events (as in ϵ_1^4 or ϵ_2^4), then ∞ is produced as the result of minimisation.

The gradient value is set to zero on source devices. For other devices, it is set by considering each neighbour ϵ' of ϵ excluding the current device itself: $N(\epsilon) \setminus p(\epsilon)$. For each such neighbour, the algorithm selects the minimum distance as computed through it, by summing up the estimates of the distance $g(\epsilon')$ from the closest source to the neighbour with the distance $D_\epsilon(d(\epsilon'))$ from the neighbour to the current event. By running such function in each event, the nodes will tend to adjust their gradient value towards the “correct” distance value. Figure 2 presents a sample execution of this function on an event structure.

In XC, the gradient function just described can be written as follows:

```
def gradient(isSource) {
  exchange(  $\infty$ , (n) => nfold(min, n + nbrDist(), mux(isSource, 0,  $\infty$ )) )
}
```

In the code above, we assume that the built-in function `nbrDist` returns in each event ϵ the nvalue $0[\bar{d} \mapsto D_\epsilon(\bar{d})]$, where $\bar{d} = N(\epsilon) \setminus p(\epsilon)$. Furthermore, the built-in function `nfold(min, \underline{w} , ℓ)` returns the minimum across ℓ and each $\underline{w}(\delta')$ for a neighbour $\delta' \neq \delta$ excluding the current device $\delta = d(\epsilon)$.

Syntax:	
$e ::= x \mid \text{fun } x(\bar{x})\{e\} \mid e(\bar{e}) \mid \text{val } x = e; e \mid \ell \mid w$	expression
$w ::= \ell[\bar{\delta} \mapsto \bar{\ell}]$	nvalue
$\ell ::= b \mid \text{fun } x(\bar{x})\{e\} \mid c(\bar{\ell})$	local literal
$b ::= \text{exchange} \mid \text{nfold} \mid \text{self} \mid \text{updateSelf} \mid \text{uid} \mid \dots$	built-in function
<hr/>	
Free variables of an expression:	
$FV(x) = \{x\} \quad FV(\ell) = FV(w) = \emptyset \quad FV(\text{fun } x_0(x_1, \dots, x_n)\{e\}) = FV(e) \setminus \{x_0, \dots, x_n\}$	
$FV(e_0(e_1, \dots, e_n)) = \bigcup_{i=0..n} FV(e_i) \quad FV(\text{val } x = e; e') = FV(e) \cup FV(e') \setminus \{x\}$	
<hr/>	
Syntactic sugar:	
$(\bar{x}) \Rightarrow e$	$::= \text{fun } y(\bar{x})\{e\}$ where y is a fresh variable
$\text{def } x(\bar{x})\{e\}$	$::= \text{val } x = \text{fun } x(\bar{x})\{e\};$
$\text{if}(e)\{e_{\top}\} \text{ else } \{e_{\perp}\}$	$::= \text{mux}(e, () \Rightarrow e_{\top}, () \Rightarrow e_{\perp})()$

Figure 3: Syntax (top), free variables (middle) and syntactic sugar (bottom) for XC expressions.

3.4. Syntax. Figure 3 (top) shows the formal syntax of XC. As in [IPW01], the overbar notation indicates a (possibly empty) sequence of elements, e.g., \bar{x} is short for x_1, \dots, x_n ($n \geq 0$). The syntax is that of a standard functional language, with no peculiar features for distribution, that are instead apparent in the semantics. An XC *expression* e can be either:

- a *variable* x ;
- a (possibly recursive) *function* $\text{fun } x(\bar{x})\{e\}$, which may have free variables;
- a *function call* $e(\bar{e})$;
- a *let-style* expression $\text{val } x = e; e$;
- a *local literal* ℓ , that is either a built-in function b , a defined function $\text{fun } x(\bar{x})\{e\}$ *without* free variables, or a data constructor c applied to local literals (possibly none);
- an *nvalue* w , as described in Section 3.3.

XC can be typed using standard higher-order let-polymorphism, as shown in the following section, without distinguishing between types for local and neighbouring values. This will reflect in the semantics by having constructs and built-in functions of the language accepting nvalues for their arguments, and by having implicit promotion of local values ℓ to nvalues $\ell[]$. Free variables are defined in a standard way (Figure 3, middle), and an expression e is *closed* if $FV(e) = \emptyset$. Programs are closed expressions without nvalues as sub-expressions: nvalues only arise in computations, and are the only values produced by evaluating programs. The minimal syntax presented is often extended through syntactic sugar to include infix operators, omitted parenthesis in 0-ary constructors, and other non-trivial encoding described in Figure 3 (bottom).

3.5. Typing. Figure 4 shows a classic Hindley-Milner type system [DM82] for XC. A type T can be:

- a *type variable* α ;
- a (recursive) *data type* $K[\bar{T}]$, consisting of a *parametric type name* K of arity $n \geq 0$, applied to n types T_1, \dots, T_n (possibly zero);
- or a *function type* $(\bar{T}) \rightarrow T$.

Types, type schemes, typing environments and instantiation:			
$T ::= \alpha \mid K[\bar{T}] \mid (\bar{T}) \rightarrow T$	type	$\mathbf{a} ::= \mathbf{x} \mid \mathbf{b} \mid \mathbf{c}$	assumption subject
$TS ::= \forall \bar{\alpha}. T$	type scheme	$\mathcal{A} ::= \bullet \mid \mathcal{A}, \mathbf{a} : TS$	typing environment
$\forall \bar{\alpha}. T \prec T[\bar{\alpha} := \bar{T}]$	instantiation		
Expression typing:			
$\frac{[\text{T-ASS}] \mathcal{A}(\mathbf{a}) \prec T \text{ for } \mathbf{a} = \mathbf{x} \text{ or } \mathbf{b}}{\mathcal{A} \vdash \mathbf{a} : T}$	$\frac{[\text{T-LIT}] \mathcal{A}(\mathbf{c}) \prec (\bar{T}) \rightarrow T \quad \mathcal{A} \vdash \ell : \bar{T}}{\mathcal{A} \vdash \mathbf{c}(\ell) : T}$	$\boxed{\mathcal{A} \vdash \mathbf{e} : T}$	
$\frac{[\text{T-NVAL}] \mathcal{A} \vdash \ell, \bar{\ell} : T \quad \bar{\delta} \text{ distinct}}{\mathcal{A} \vdash \ell[\bar{\delta} \mapsto \bar{\ell}] : T}$	$\frac{[\text{T-APP}] \mathcal{A} \vdash \mathbf{e} : (\bar{T}) \rightarrow T \quad \mathcal{A} \vdash \bar{\mathbf{e}} : \bar{T}}{\mathcal{A} \vdash \mathbf{e}(\bar{\mathbf{e}}) : T}$		
$\frac{[\text{T-FUN}] \mathcal{A}, \mathbf{x} : \forall \bullet. (\bar{T}) \rightarrow T, \bar{\mathbf{x}} : \forall \bullet. \bar{T} \vdash \mathbf{e} : T}{\mathcal{A} \vdash \mathbf{fun} \mathbf{x}(\bar{\mathbf{x}})\{\mathbf{e}\} : (\bar{T}) \rightarrow T}$			
$\frac{[\text{T-VAL}] \mathcal{A} \vdash \mathbf{e}_1 : T_1 \quad \bar{\alpha} = \text{TV}(T_1) \quad \mathcal{A}, \mathbf{x} : \forall \bar{\alpha}. T_1 \vdash \mathbf{e}_2 : T_2}{\mathcal{A} \vdash \mathbf{val} \mathbf{x} = \mathbf{e}_1; \mathbf{e}_2 : T_2}$			

Figure 4: Typing of XC expressions.

Given that local values and nvalues are treated uniformly in the semantics of XC, and only nvalues result from evaluating closed expressions, no special types are needed for nvalues: each nvalue takes on the (classic) type of its messages. The underlining notation for nvalue types and variables used before is simply for readability and not part of the type system itself. The set of type variables in T is denoted by $\text{TV}(T)$. Polymorphism in functions and data constructors is supported by type schemes TS of the form $\forall \bar{\alpha}. T$, where $\bar{\alpha}$ are free in T , representing all types obtained by substituting $\bar{\alpha}$ with types \bar{T} , as described by the type scheme instantiation relation \prec .

A *typing environment* \mathcal{A} is a set of *assumptions* $\mathbf{a} : TS$ where the \mathbf{a} can be a variable, built-in function or constructor. Sensor types \mathbf{s} are assumed to be of the form $() \rightarrow T$ with $\text{TV}(T) = \emptyset$. For program typing, an *initial typing environment* \mathcal{A}_0 defines a type scheme for each available constructor and built-in function. This environment is extended with further assumptions for bounded variables as they appear within sub-expressions of the program. We denote by $\mathcal{A}(\mathbf{a})$ the unique type scheme assigned to \mathbf{a} in \mathcal{A} .

Expression typing is specified by judgement $\mathcal{A} \vdash \mathbf{e} : T$ meaning “expression \mathbf{e} has type T under assumptions \mathcal{A} ”. Following [IPW01], we expand multiple overbars together, so for instance $\mathcal{A} \vdash \bar{\mathbf{e}} : \bar{T}$ represents $\mathcal{A} \vdash \mathbf{e}_1 : T_1, \dots, \mathcal{A} \vdash \mathbf{e}_n : T_n$. Typing rules are all syntax-directed and mostly standard, except for [T-NVAL] that is XC-specific. This rule ensures that all messages $\ell, \bar{\ell}$ within an nvalue share the same type T and that $\bar{\delta}$ has no duplicate entries.

3.6. Semantics. The XC semantics is defined as (i) a big-step operational *device semantics*, providing a formal account of the computation of a device within one round; and (ii) a denotational *network semantics*, formalising how different device rounds communicate.

Device semantics. Figure 5 presents the device semantics, formalised by judgement $\delta; \sigma; \Theta \vdash \mathbf{e} \Downarrow \mathbf{w}; \theta$, to be read as “expression \mathbf{e} evaluates to nvalue \mathbf{w} and value-tree θ on device δ with respect to sensor values σ and value-tree environment Θ ”, where:

Auxiliary definitions:			
$\theta ::= \langle \bar{\theta} \rangle \mid \mathbf{w} \langle \bar{\theta} \rangle$	value-tree	σ	sensor state
$\Theta ::= \bar{\delta} \mapsto \bar{\theta}$	value-tree environment	δ	device identifier
$\pi_i(\langle \theta_1, \dots, \theta_n \rangle) = \theta_i$	$\pi_i(\mathbf{w} \langle \theta_1, \dots, \theta_n \rangle) = \theta_i$	$\pi_i(\bar{\delta} \mapsto \bar{\theta}) = \bar{\delta} \mapsto \pi_i(\bar{\theta})$	
$\bar{\delta} \mapsto \bar{\theta} \upharpoonright_{\mathbf{f}} = \left\{ \delta_i \mapsto \theta_i \mid \theta_i = \mathbf{w} \langle \bar{\theta}' \rangle, \text{name}(\mathbf{w}(\delta_i)) = \text{name}(\mathbf{f}) \right\}$			
		$\text{name}(\mathbf{b}) = \mathbf{b}$	$\text{name}(\mathbf{fun}^\tau \mathbf{x}(\bar{\mathbf{x}})\{\mathbf{e}\}) = \tau$
Evaluation rules:			
$\delta; \sigma; \Theta \vdash \mathbf{e} \Downarrow \mathbf{w}; \theta$			
$\frac{[\text{E-NBR}]}{\delta; \sigma; \Theta \vdash \mathbf{w} \Downarrow \mathbf{w}; \langle \rangle}$	$\frac{[\text{E-VAL}]}{\delta; \sigma; \pi_1(\Theta) \vdash \mathbf{e}_1 \Downarrow \mathbf{w}_1; \theta_1 \quad \delta; \sigma; \pi_2(\Theta) \vdash \mathbf{e}_2[\mathbf{x} := \mathbf{w}_1] \Downarrow \mathbf{w}_2; \theta_2}{\delta; \sigma; \Theta \vdash \mathbf{val} \mathbf{x} = \mathbf{e}_1; \mathbf{e}_2 \Downarrow \mathbf{w}_2; \langle \theta_1, \theta_2 \rangle}$		
$\frac{[\text{E-LIT}]}{\delta; \sigma; \Theta \vdash \ell \Downarrow \ell[]; \langle \rangle}$			
$\frac{[\text{E-APP}]}{\delta; \sigma; \Theta \vdash \mathbf{e}_0(\mathbf{e}_1, \dots, \mathbf{e}_n) \Downarrow \mathbf{w}_{n+1}; \mathbf{f}[] \langle \theta_0, \dots, \theta_{n+1} \rangle}$			
$\delta; \sigma; \pi_{i+1}(\Theta) \vdash \mathbf{e}_i \Downarrow \mathbf{w}_i; \theta_i \quad \text{for all } i \in 0, \dots, n$			
$\delta; \sigma; \pi_{n+2}(\Theta \upharpoonright_{\mathbf{f}}) \vdash \mathbf{f}(\mathbf{w}_1, \dots, \mathbf{w}_n) \Downarrow^* \mathbf{w}_{n+1}; \theta_{n+1} \text{ where } \mathbf{f} = \mathbf{w}_0(\delta)$			
Auxiliary evaluation rules:			
$\delta; \sigma; \Theta \vdash \mathbf{f}(\bar{\mathbf{w}}) \Downarrow^* \mathbf{w}; \theta$			
$\frac{[\text{A-FUN}]}{\delta; \sigma; \Theta \vdash \mathbf{e}[\mathbf{x} := \mathbf{fun}^\tau \mathbf{x}(\bar{\mathbf{x}})\{\mathbf{e}\}, \bar{\mathbf{x}} := \bar{\mathbf{w}}] \Downarrow \mathbf{w}; \theta \quad \delta; \sigma; \Theta \vdash \mathbf{fun}^\tau \mathbf{x}(\bar{\mathbf{x}})\{\mathbf{e}\}(\bar{\mathbf{w}}) \Downarrow^* \mathbf{w}; \theta}$	$\frac{[\text{A-UID}]}{\delta; \sigma; \Theta \vdash \mathbf{uid}() \Downarrow^* \delta; \langle \rangle}$		
$\Theta = \bar{\delta} \mapsto \bar{\mathbf{w}} \langle \dots \rangle \quad \mathbf{w} = \mathbf{w}_i[\bar{\delta} \mapsto \bar{\mathbf{w}}(\delta)] \quad \mathbf{f} = \mathbf{w}_f(\delta)$			
$\frac{[\text{A-XC}]}{\delta; \sigma; \Theta \vdash \mathbf{exchange}(\mathbf{w}_i, \mathbf{w}_f) \Downarrow^* \mathbf{w}_r; \mathbf{w}_s \langle \theta \rangle}$	$\frac{[\text{A-SELF}]}{\delta; \sigma; \Theta \vdash \mathbf{self}(\mathbf{w}) \Downarrow^* \mathbf{w}(\delta); \langle \rangle}$		
$\Theta = \delta_1 \mapsto \theta_1, \dots, \delta_n \mapsto \theta_n \quad \ell_0 = \mathbf{w}_3(\delta) \quad \mathbf{f} = \mathbf{w}_1(\delta)$			
$\frac{[\text{A-FOLD}]}{\delta; \sigma; \Theta \vdash \mathbf{nfold}(\mathbf{w}_1, \mathbf{w}_2, \mathbf{w}_3) \Downarrow^* \ell_n[]; \langle \rangle}$			
$\delta; \sigma; \emptyset \vdash \mathbf{f}(\ell_{i-1}[], \mathbf{w}_2(\delta_i)[]) \Downarrow^* \ell_i[]; \theta \text{ if } \delta_i \neq \delta \text{ else } \ell_i = \ell_{i-1} \quad \dots$			

Figure 5: Device big-step operational semantics of XC.

- \mathbf{w} is called the *result* of \mathbf{e} ;
- θ is an ordered tree with nvalues on some nodes (cf. Figure 5 (top)), representing messages to be sent to neighbours by tracking the nvalues produced by exchange-expressions in \mathbf{e} , and the stack frames of function calls;
- Θ collects the (non expired) value-trees received by the most recent firings of neighbours of δ , as a map $\delta_1 \mapsto \theta_1, \dots, \delta_n \mapsto \theta_n$ ($n \geq 0$) from device identifiers to value-trees.

The semantics of function application is expressed through an auxiliary evaluation judgement of similar form, $\delta; \sigma; \Theta \vdash \mathbf{f}(\bar{\mathbf{w}}) \Downarrow^* \mathbf{w}; \theta$, to be read as “expression $\mathbf{f}(\bar{\mathbf{w}})$ evaluates to nvalue \mathbf{w} and value-tree θ on device δ with respect to sensor values σ and value-tree environment Θ ”. We assume every function expression $\mathbf{fun} \mathbf{x}(\bar{\mathbf{x}})\{\mathbf{e}\}$ occurring in the program is annotated with a unique name τ before the evaluation starts. Then, τ will be the name for the annotated function expression $\mathbf{fun}^\tau \mathbf{x}(\bar{\mathbf{x}})\{\mathbf{e}\}$, and \mathbf{b} the name for a built-in function \mathbf{b} .

The syntax of value-trees and value-tree environments is in Figure 5 (top). The rules for judgement $\delta; \sigma; \Theta \vdash \mathbf{e} \Downarrow \mathbf{v}; \theta$ (Figure 5, middle) are standard for functional languages, extended to evaluate a sub-expression \mathbf{e}' of \mathbf{e} w.r.t. the value-tree environment Θ' obtained from Θ by extracting the corresponding subtree (when present) in the value-trees in the range of Θ . This *alignment* process is modelled by the auxiliary “projection” functions π_i (for

any positive natural number i) (Figure 5, top). When applied to a value-tree θ , π_i returns the i -th sub-tree θ_i of θ . When applied to a value-tree environment Θ , π_i acts pointwise on the value-trees in Θ . The alignment process ensures that the value-trees in the environment Θ always correspond to the evaluation of the same sub-expression currently being evaluated. To ensure this match holds (as said before, of the stack frame and position in the AST), in the evaluation of a function application $\mathbf{f}(\bar{\mathbf{w}})$, the environment Θ is reduced to the smaller set $\Theta|_{\mathbf{f}}$ of trees which corresponded to the evaluation of a function with the same *name*.

Rule [E-NBR] evaluates an nvalue \mathbf{w} to \mathbf{w} itself and the empty value-tree. Rule [E-LIT] evaluates a local literal ℓ to the nvalue $\ell[]$ and the empty value-tree. Rule [E-VAL] evaluates a val-expression, by evaluating the first sub-expression with respect to the first sub-tree of the environment obtaining a result \mathbf{w}_1 , and then the second sub-expression with respect to the second sub-tree of the environment, after substituting the variable \mathbf{x} with \mathbf{w}_1 .

Rule [E-APP] is standard eager function application: the function expression \mathbf{e}_0 and each argument \mathbf{e}_i are evaluated w.r.t. $\pi_{i+1}(\Theta)$ producing result \mathbf{v}_i and value-tree θ_i . Then, the function application itself is demanded to the auxiliary evaluation rules, w.r.t. the last sub-tree of the trees corresponding to the same function: $\pi_{n+2}(\Theta|_{\mathbf{f}})$. The auxiliary rule [A-FUN] handles the application of fun-expression, which evaluates the body after replacing the arguments $\bar{\mathbf{x}}$ with their provided values $\bar{\mathbf{w}}$, and the function name \mathbf{x} with the fun-expression itself to allow for recursion. Rules [A-UID] and [A-SELF] trivially encode the behaviour of the **uid** and **self** built-ins. Rule [A-XC] evaluates an exchange-expression, realising the behaviour described at the end of Section 3.3. Notation $\mathbf{w}_1[\bar{\delta} \mapsto \bar{\mathbf{w}}(\delta)]$ is used to represent the nvalue \mathbf{w}_1 after the update for each i of the message for δ_i with the custom message $\mathbf{w}_i(\delta)$. The result is fed as argument to function \mathbf{w}_f : the first element of the resulting pair is the overall result, while the second is used to tag the root of the value-tree. Rule [A-FOLD] encodes the **nfold** operators. First, the domain of Θ is inspected, giving a (sorted) list $\delta_1, \dots, \delta_n$. An initial local value ℓ_0 is set to the “self-message” of the third argument. Then, a sequence of ℓ_i is defined, each by applying function \mathbf{w}_1 to the previous element in the sequence and the value $\mathbf{w}_2(\delta_i)$ (skipping δ itself). The final result ℓ_n is the result of the application, with empty value-tree. Auxiliary rules for the other available built-in functions are standard, do not depend on the environment, hence have been omitted.

Example 3.2 (Gradient Semantics). Let us consider the semantics of the gradient function presented in Example 3.1, as used in the following program:

```
def gradient(isSource) {
  exchange( ∞, (n) => nfold(min, n + nbrDist(), mux(isSource, 0, ∞)) )
}
gradient(sourceSensor())
```

After removing all syntactic sugar (described in Figure 3), the program converts to:

```
val gradient = fun gradient(isSource) {
  exchange( ∞, fun f(n) { nfold(min, n + nbrDist(), mux(isSource, 0, ∞)) } )
};
gradient(sourceSensor())
```

Consider the execution of such program in a given device δ , with respect to sensor values σ and value-tree environment Θ . The overall program is a val-expression, hence rule [E-VAL] applies. This rule demands the evaluation of the first sub-expression first (w.r.t. $\pi_1(\Theta)$), which is **fun gradient**... This is a local literal defining a function, so rule [E-LIT] applies, and produces an nvalue without exceptions that corresponds to the function (and empty

value-tree). Then, this resulting value is substituted for the variable `gradient` in the last part of the program, obtaining:

```
fun gradient(isSource) {
  exchange( ∞, fun f(n) { nfold(min, n + nbrDist(), mux(isSource, 0, ∞)) } )
}[] (sourceSensor())
```

According to rule [E-VAL], this code has to be evaluated (w.r.t $\pi_2(\Theta)$). This is a function application, so rule [E-APP] applies. This evaluates the function first (w.r.t. $\pi_1(\pi_2(\Theta))$), which is already an nvalue, so it evaluates to itself (and no value-tree) through rule [E-NBR]. Then, it evaluates the argument `sourceSensor()` (w.r.t. $\pi_2(\pi_2(\Theta))$), which is a further function application. Without describing this further evaluation in details for brevity, it will ultimately produce a Boolean value b that is extracted from the sensor status σ through the auxiliary evaluation rule defined for the built-in function `sourceSensor`.

As the last step of the overall function application, the auxiliary predicate is used, applying rule [A-FUN] to expand the function body (w.r.t. $\pi_3(\pi_2(\Theta))$). The substitution results in the following code to be evaluated:

```
exchange( ∞, fun f(n) { nfold(min, n + nbrDist(), mux(b[], 0, ∞)) } )
```

This is another function application where all expressions are local literals, so it resorts to rule [E-APP], where the only interesting part is evaluating the application itself (w.r.t. $\Theta' = \pi_4(\pi_3(\pi_2(\Theta)))$) with the auxiliary rule [A-XC]. If we consider δ, σ, Θ to be as in event ϵ_5^2 described in Figure 2, we have that $b = \perp$, $\delta = \delta_2$ and

$$\Theta' = \delta_1 \mapsto 4[] \langle \dots \rangle, \delta_2 \mapsto \infty[] \langle \dots \rangle, \delta_3 \mapsto 1[] \langle \dots \rangle.$$

According to rule [A-XC], first we compute a nvalue combining received messages (using the first argument as default), which in this case is $w = \infty[\delta_1 \mapsto 4, \delta_2 \mapsto \infty, \delta_3 \mapsto 1]$. This value is fed to the second argument, resorting in the function application:

```
fun f(n) { nfold(min, n + nbrDist(), mux(⊥[], 0, ∞)) } (∞[δ1 ↦ 4, δ2 ↦ ∞, δ3 ↦ 1])
```

Which after a further application of [A-FUN] translates to the evaluation of:

```
nfold(min, ∞[δ1 ↦ 4, δ2 ↦ ∞, δ3 ↦ 1] + nbrDist(), mux(⊥[], 0, ∞))
```

The evaluation proceeds similarly. Following the example in Figure 2, `nbrDist()` evaluates to $1[]$, which added to w becomes $\infty[\delta_1 \mapsto 5, \delta_2 \mapsto \infty, \delta_3 \mapsto 2]$, and the `mux` call evaluates to $\infty[]$, leading to:

```
nfold(min, ∞[δ1 ↦ 5, δ2 ↦ ∞, δ3 ↦ 2], ∞[])
```

According to rule [A-FOLD], we set $l_0 = \infty$, $l_1 = \min(l_0, 5) = 5$, $l_2 = l_1$, $l_3 = \min(l_2, 2) = 2$, which is the overall result (and is stored in the value-tree as well by rule [A-XC]).

Network semantics. The evolution of a network executing a program is formalised through structures of *events*, atomic computations performed according to the device-level semantics. Such events across space and time (i.e., when and in which devices they happen) form an “aggregate” execution in which information is exchanged across events following a *messaging* relationship \rightsquigarrow . This idea, informally presented in Section 3.2, can be formalised through *augmented event structures*, which expands a classical *event structure* [Lam78] with further information associated to events: device identifiers and sensor status.

Definition 3.3 (Augmented Event Structure). An *augmented event structure* $\mathbb{E} = \langle E, \rightsquigarrow, d, s \rangle$ is a tuple where:

- E is a countable set of *events* ϵ ,
 - $\rightsquigarrow \subseteq E \times E$ is a *messaging* relation,
 - $d : E \rightarrow \Delta$ is a mapping from events to the devices where they happened, and
 - $s : E \rightarrow S$ maps each event ϵ to a *sensors status* σ (as in the device-level semantics),
- such that:

- $\epsilon_1 = \epsilon_2$ whenever $d(\epsilon_1) = d(\epsilon_2)$ and $\epsilon_i \rightsquigarrow \epsilon$ for $i = 1, 2$ (i.e., all predecessors of an event happened on different devices);
- there are no sequences $\epsilon_1 \rightsquigarrow \dots \rightsquigarrow \epsilon_n \rightsquigarrow \epsilon_1$ (i.e., the \rightsquigarrow relation is acyclic);
- the set $X_\epsilon = \{\epsilon' \in E \mid \epsilon' \rightsquigarrow \dots \rightsquigarrow \epsilon\}$ of events that can reach ϵ in \rightsquigarrow is finite for all ϵ (i.e., \rightsquigarrow is well-founded and locally finite).

We say that event ϵ' is a *neighbour* of event ϵ iff $\epsilon' \rightsquigarrow \epsilon$. We call *causality relation* the irreflexive partial order $< \subseteq E \times E$ obtained as the transitive closure of \rightsquigarrow .

This computational model captures the behaviour of communicating physical devices. The evolution of the state of devices is a chain of \rightsquigarrow -connected events associated to the same device. Communication is captured by \rightsquigarrow -connected events associated to different devices. Such a computational model allows us to express programs abstracting from synchronisation, shared clocks, or regularity and frequency of events. Informally, following previous work [MZ04, LLM17, VAB⁺18], we refer to a *field of values* as a mapping from devices to values, capturing a global snapshot of the values produced by the most recent firing of each device. The evolution over time of a field of values is then a *space-time value*, which maps each event (in an augmented event structure) to a corresponding value.

Definition 3.4 (Space-Time Value). Let $\mathbb{E} = \langle E, \rightsquigarrow, d, s \rangle$ be an augmented event structure and \mathbf{V}_T be the domain of nvalues of type T . A *space-time value* (in \mathbb{E} of type T) is a function mapping events to nvalues $\Phi : E \rightarrow \mathbf{V}_T$.

In this model, the evaluation of a program e_{main} in an augmented event structure \mathbb{E} then produces a space-time value, induced by repeatedly applying the device-level semantics.

Definition 3.5 (Program Evaluation on Event structures). Let e be an XC expression of type T given assumptions \mathcal{A} . Let \mathbb{E} be an augmented event structure whose s includes values of the appropriate type for each sensor and free variable appearing in e .

Let $\theta_e^{\mathbb{E}} : E \rightarrow \Omega$ (where Ω is the set of all value-trees) and $\Phi_e^{\mathbb{E}} : E \rightarrow \mathbf{V}_T$ be defined by induction on ϵ in E , so that $d(\epsilon); s(\epsilon); \Theta_\epsilon \vdash e \Downarrow \Phi_e^{\mathbb{E}}(\epsilon); \theta_e^{\mathbb{E}}(\epsilon)$ where $\Theta_\epsilon = \{d(\epsilon') \mapsto \theta_e^{\mathbb{E}}(\epsilon') \mid \epsilon' \rightsquigarrow \epsilon\}$. Then we say that $\Phi_e^{\mathbb{E}} : E \rightarrow \mathbf{V}_T$ is the evaluation of expression e on \mathbb{E} .

4. DISTRIBUTED COLLECTIVE PROCESSES IN XC

In this section, we present an implementation of DCPs in XC. First, we characterise the implementation in terms of an abstract notation and formulas on event structures (Section 4.1). Then, we provide a formalisation of DCPs in terms of the big-step operational semantics for a new XC construct (Section 4.2), that can be used to actually *program* DCPs. We conclude discussing engineering characteristics of DCPs (Section 4.3) and their practical implementation in the FCPP framework (Section 4.4).

4.1. Modelling on event structures. A *distributed collective process (DCP)* P is a computation with given programmed behaviour. A single DCP can be run in multiple *process instances* P_i , each associated to a unique *process identifier (PID)* i , which we assume also embeds construction parameters for the process instance. New instances of an aggregate process P are spawned through a *generation field* G_P , producing a set of identifiers $G_P(\epsilon) = \{i \dots\}$ in each event ϵ , of process instances that need to be created in that event ϵ (which we call *initiator* for P_i). For each process instance P_i , we use the Boolean predicate $\pi_{P_i}(\epsilon)$ to denote whether such instance is being executed at event ϵ (either being initiated by ϵ , or through propagation from previous events). Each process instance P_i , if active in an event ϵ (i.e., $\pi_{P_i}(\epsilon) = \top$), locally computes both an *output* $O_{P_i}(\epsilon)$ (returned to the process caller) and a *status* $s_{P_i}(\epsilon)$ according to some given program (possibly involving communication with neighbours). The status is an nvalue mapping the device $d(\epsilon')$ of each neighbour event $\epsilon' \in N$ to a `bool` value indicating whether that device has to run the same process instance the next round. The default value of the status determines whether the process instance should be propagated or not to devices that were not neighbours in ϵ . Every event that has *at least one* neighbour that propagated the process instance to it will run the process instance as well. More precisely, a process instance P_i which is active in an event ϵ *potentially* propagates the process to any event ϵ' of which ϵ for which it is a neighbour ($\epsilon \rightsquigarrow \epsilon'$) depending on the value of $s_{P_i}(\epsilon)$. In formulas:

$$\pi_{P_i}(\epsilon) = \begin{cases} \top & \text{if } i \in G_P(\epsilon), \text{ otherwise} \\ \top & \text{if } \exists \epsilon' \rightsquigarrow \epsilon. \pi_{P_i}(\epsilon') \wedge s_{P_i}(\epsilon')(d(\epsilon)) = \top \\ \perp & \text{otherwise.} \end{cases}$$

This notion of DCPs, defined by the abstract model above, is particularly well-suited for being implemented in XC, as shown in the following section.

Example 4.1 (Multi-gradient). As a running example, consider a *multi-gradient* P^{mg} : the construction of distinct gradients (cf. Example 3.1) from multiple sources. A gradient has to be spawned locally in source devices; therefore, the generation logic $G_{P^{mg}}$ could be as follows:

$$G_{P^{mg}}(\epsilon) = \begin{cases} \{d(\epsilon)\} & \text{isSource}(\epsilon) \\ \emptyset & \text{otherwise} \end{cases}$$

where $\text{isSource}(\epsilon)$ is true iff the device $d(\epsilon)$ is considered a source when event ϵ occurs. The output $O_{P_i^{mg}}(\epsilon)$ is the gradient as computed in ϵ , representing the distance from source device i to the current device $d(\epsilon)$. This can be modelled by the following output computation, analogous to the computation of $g(\epsilon)$ as presented in Example 3.1:

$$O_{P_i^{mg}}(\epsilon) = \begin{cases} \min_{\epsilon' \in N(\epsilon) \setminus p(\epsilon)} \{O_{P_i^{mg}}(\epsilon') + D_\epsilon(d(\epsilon'))\} & d(\epsilon) \neq i \\ 0 & d(\epsilon) = i \end{cases}$$

If each gradient has to be limited in space, say within a maximum distance of ρ from the source, then a possible implementation of the status logic $s_{P_i^{mg}}$ of a process instance could be the following:

$$s_{P_i^{mg}}(\epsilon) = \left(O_{P_i^{mg}}(\epsilon) \leq \rho \right) [i \mapsto \perp]$$

In this way, $s_{P_i^{mg}}(\epsilon)$ is an nvalue that by default correspond to the Boolean value of whether the current gradient value is less than or equal to ρ . This nvalue carries a single exception,

Auxiliary definitions:	
$\theta ::= \langle \bar{\theta} \rangle \mid \mathbf{w}(\bar{\theta}) \mid \bar{\ell} \mapsto \bar{\theta}$	value-tree
$\pi^\ell(\bar{\ell} \mapsto \bar{\theta}) = \theta_i$	s.t. $\ell_i = \ell$ if it exists else \bullet
Auxiliary evaluation rules:	
$\delta; \sigma; \Theta \vdash \mathbf{f}(\bar{\mathbf{w}}) \Downarrow^* \mathbf{w}; \theta$	
[A-SPAWN]	$\frac{k_1, \dots, k_n = \mathbf{w}^k(\delta) \cup \{k \text{ for } \delta' \in \mathbf{dom}(\Theta), k \mapsto b(\theta) \in \Theta(\delta') \text{ with } b(\delta) = \mathbf{True}\}}{\delta; \sigma; \pi_1(\pi^{k_i}(\Theta)) \vdash \mathbf{w}^p(k_i) \Downarrow \mathbf{w}_i; \theta_i \text{ where } \mathbf{w}_i = \mathbf{pair}(\mathbf{v}_i, b_i) \text{ for } i \in 1, \dots, n}$ $\delta; \sigma; \Theta \vdash \mathbf{spawn}(\mathbf{w}^p, \mathbf{w}^k) \Downarrow^* \bar{k} \mapsto \bar{\mathbf{w}}; \bar{k} \mapsto \bar{b}(\bar{\theta})$

Figure 6: Device (big-step) operational semantics of XC

for the source device i , for which the value \perp is used instead, so that the process instance is never propagated to the source. This makes sure that a the process instance can live on the source only as long as the source is in the generation set G_{Pmg} (and thus *isSource* is still true for it). In case i stops being a source, the source itself would immediately leave the process, and for other devices the gradient estimates in $O_{P_i^{mg}}(\epsilon)$ would start to rise, until they surpass the threshold ρ effectively terminating the whole instance.

4.2. Formalisation in XC. In XC, the notion of DCPs is implemented through a single built-in construct $spawn_{XC}(P, G_P)$ that runs independent instances of a given process P , where new instances are locally generated according to *generation field* G_P as explained above. The output of a $spawn_{XC}(P, G_P)$ expression in an event ϵ is the set of pairs $\{(i, O_{P_i}(\epsilon)), \dots\}$ for which $\pi_{P_i}(\epsilon) = \top$. See Figure 7 for an example showing two overlapping process instances.

The $spawn_{XC}$ construct corresponds to the mathematical definition given in the previous section (Section 4.1), and embeds naturally in XC as a built-in function. As a built-in in XC, **spawn** assumes the same type as the classical spawn construct in Field Calculus (FC): $\forall \alpha_k, \alpha_v. ((\alpha_k) \rightarrow \mathbf{pair}[\alpha_v, \mathbf{bool}], \mathbf{set}[\alpha_k]) \rightarrow \mathbf{map}[\alpha_k, \alpha_v]$. However, in XC every type allows nvalues, which translates into practical differences.

Figure 6 presents the semantics of the spawn built-in, relative to the XC semantics presented in Section 3.6. In order to introduce the **spawn** construct, it is necessary to extend the auxiliary definition of value-trees (highlighted in grey), to also allow for maps from local literals ℓ (identifiers of the running processes) to their corresponding value-trees. Then, rule [A-SPAWN] can be written by naturally porting the similar rule in [CVA⁺19], while using the fact that the Boolean returned by the process is an nvalue, and thus can be different for different neighbours. In this rule, a list of *process keys* \bar{k} is computed by adjoining (i) the keys $\mathbf{w}^k(\delta)$ currently present in the second argument \mathbf{w}^k of **spawn** for the current device δ ; (ii) the keys that any neighbour δ' broadcast in their last message $\Theta(\delta')$, provided that the corresponding Boolean value b returned was true for the current device $b(\delta) = \mathbf{True}$ (thus, demanding process expansion to δ). To realise “multiple alignment”, for each key k_i , the process \mathbf{w}^p is applied to k_i with respect to the part of the value-tree environment $\pi_1(\pi^{k_i}(\Theta))$ that corresponds to key k_i , producing $\mathbf{w}_i; \theta_i$ as a result. Finally, the construct concludes returning the maps $\bar{k} \mapsto \bar{\mathbf{w}}; \bar{k} \mapsto \bar{b}(\bar{\theta})$ mapping process keys to their evaluation result.

Example 4.2 (Multi-gradient Implementation). As an example, the multi-gradient introduced in Example 4.1 can be translated in XC obtaining the following code:

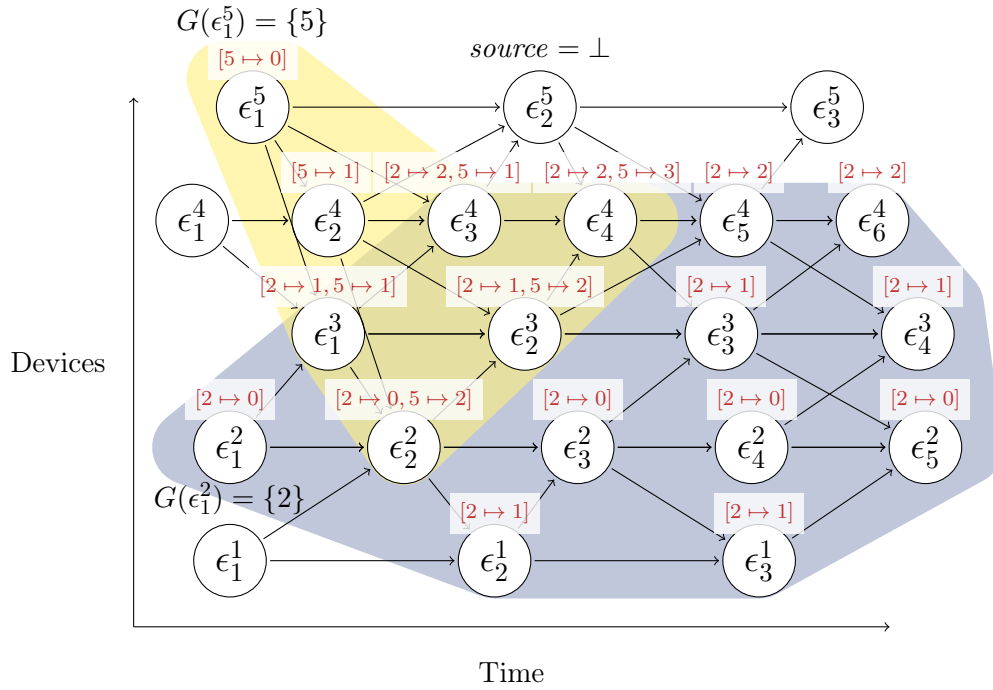


Figure 7: This diagram shows the dynamics of a multi-gradient (cf. Example 4.1) with devices 2 and 5 as sources, where the former is a source for the entire execution shown, and the latter only in its first event, ϵ_1^5 . Two process instances P_5 (yellow) and P_2 (blue), each one carrying a gradient process, are generated from events ϵ_1^5 and ϵ_1^2 , respectively. Suppose both processes are instantiated to propagate within a maximum number $\rho = 1$ of hops from the initiator device. The two process instances *overlap* in various events of devices 2 (ϵ_2^2 , ϵ_3^2), 3 (ϵ_3^3 , ϵ_2^3), and 4 (ϵ_4^3 , ϵ_4^4)—where both instances are run at the same time in non-deterministic order. Since it is assumed that the device 5 is a source only at ϵ_1^5 ; then, devices 4, 3, and 2 start bubbling the gradient value up, until exiting the process. Recall that any event propagates the process if and only if its gradient value is less than or equal $\rho = 1$, that process instances are independent, and that the gradient values adds a unit of distance to the minimum across neighbours events (excluding self).

```

def nbr(v) {
  exchange( v, (n) => (n,v) )
}
def multi_gradient(isSource, theta) {
  val gen_set = if (isSource) {set(oid())} else {set()};
  spawn((i) => {
    val output = gradient(oid() == i);
    val status = output <= theta && nbr(oid()) != i;
    (output, status)
  }, gen_set);
}

```

The multi-gradient function receives as a parameter the Boolean value *isSource* identifying source devices and the parameter ρ regulating the size of the areas spanned by each process.

The generation set `gen_set` is computed following exactly the definition of G_{Pmg} , and similarly for `output` with $O_{P_i^{mg}}$, by relying on the XC implementation of the gradient discussed in Example 3.1. In order to define `status` as in $s_{P_i^{mg}}$, the translation starts from the same default value `output <= theta` of $s_{P_i^{mg}}$, and adds the exception for i by conjunction with an nvalue that is false only for i : `nbr(uid()) != i`.

In the evaluation of the `multi_gradient` function, the application of `spawn` is performed according to rule [A-SPAWN]. First, the set of currently active process keys is retrieved from the value-tree environment Θ . For instance, in the example described in Figure 7, for event ϵ_2^3 , we would have $isSource = \perp$, $\rho = 1$, and Θ equal to:

$$\begin{aligned} \delta_2 &\mapsto (\delta_2 \mapsto \top[\delta_2 \mapsto \perp]\langle\theta_2^2\rangle, \delta_5 \mapsto \perp[]\langle\theta_2^5\rangle), \\ \delta_3 &\mapsto (\delta_2 \mapsto \top[\delta_2 \mapsto \perp]\langle\theta_3^2\rangle, \delta_5 \mapsto \top[\delta_5 \mapsto \perp]\langle\theta_3^5\rangle), \\ \delta_4 &\mapsto (\delta_5 \mapsto \top[\delta_5 \mapsto \perp]\langle\theta_4^5\rangle) \end{aligned}$$

In this case, the generation set is empty ($w^k(\delta)$ in rule [A-SPAWN]). Thus, the set of active process keys is computed only by considering, for every δ' in $\mathbf{dom}(\Theta)$ (that is, $\delta_2, \delta_3, \delta_4$) their corresponding keys *only if* associated with a true Boolean for the current device δ_3 . The key δ_5 coming from the neighbour δ is discarded, since its corresponding status is $\perp[]$ not true for δ_3 . All other keys have a corresponding status that is true for δ_3 , so the keys computed are $k_1 = \delta_2, k_2 = \delta_5$.

Then, for $i = 1, 2$ the process body is executed with respect to the environment corresponding to the same process execution in neighbours. For example, for $k_1 = \delta_2$:

$$\pi^{\delta_2}(\Theta) = \delta_2 \mapsto \top[\delta_2 \mapsto \perp]\langle\theta_2^2\rangle, \delta_3 \mapsto \top[\delta_2 \mapsto \perp]\langle\theta_3^2\rangle, \quad \pi_1(\pi^{\delta_2}(\Theta)) = \delta_2 \mapsto \theta_2^2, \delta_3 \mapsto \theta_3^2$$

and the body of the process is executed with argument δ_2 w.r.t. this environment. The execution will result in the pair $(1[], \top[\delta_2 \mapsto \perp])$ which will contribute by adding $\delta_2 \mapsto 1[]$ to the overall result and $\delta_2 \mapsto \top[\delta_2 \mapsto \perp]\langle\dots\rangle$ to the produced value-tree.

4.3. The DCP abstraction. The crucial problem that we investigate in this paper revolves around the definition of collaborative activities carried out by *dynamic* collections of devices (a.k.a. *ensembles* [NLPT14]). We call these *DCPs* since they are defined by a common control program that regulates the behaviour of a largely homogeneous set of devices.

In particular, a device may participate *concurrently* to multiple collective processes, or to multiple ensembles; and in any single round, a device executes the computation associated to *all* the currently joined DCPs. Participation to multiple DCPs may thus have consequences on local resource usage (cf. resource-constrained devices). Therefore, real-world implementations have also to take performance issues into account. Since there is no universal solution for managing resource constraints in devices participating to multiple DCPs, much of the responsibility on handling these situations lies with the programmer. Some general strategies may be applicable, such as: (i) tuning the scheduling policy according to the number of processes joined by a device, using extensions of the calculus that have been previously studied [PCV⁺21] and are already available in current implementations (including FCPP [Aud20]); (ii) computing a status from the resource availability perceived through a sensor or runtime checks, using it to prevent process propagation to nodes that are already on a high load; (iii) designing DCPs in order to minimise resource consumption, both through low-level optimisations and by reducing the area spanned by a single DCPs to the smallest possible. Further details on these strategies go beyond the scope of this paper.

We define as the *domain* of a DCP the set of the nodes that are currently running it. We define as the *shape* of a DCP the spatiotemporal region that is identified by the spatiotemporal locations of the nodes belonging to the domain of the DCP. Often, DCPs are *transient*, i.e., they have a limited lifetime: they start to exist at some time, and they dissolve once no more nodes run them.

In this work, we are mainly concerned with studying how to create and manipulate these DCPs. The supporting formal framework and implementation is described in Section 4. In summary, the developer has the following mechanisms for defining systems of DCPs:

- *generation logic*: the need for collective activities can be encoded in a rule for generating new instances of DCPs;
- *identity logic*: the logic used to identify DCPs can be used to distinguish between them and hence to regulate their domains (e.g., for controlling the granularity of teams);
- *internal logic*: this logic defines a collective computation (scoped within the domain of a single DCP) promoting decentralised decision-making, e.g., in terms of typical self-organisation patterns (collection, propagation, leader election, evaporation, etc.);
- *shape control logic*: it is possible to specify rules for the local expansion of the domains of DCPs (e.g., to gather more participants), typically also leveraging results from the internal computation itself—the XC implementation provides an especially flexible way to specify this, as different neighbours can receive different information;
- *termination logic*: this logic, strictly related to shape control, enables to specify how individual agents may leave a DCP instance as well as how an entire DCP may be terminated;
- *input logic*: existing DCPs may also be controlled by specifying “external inputs” provided as explicit arguments or closed over a lambda closure—an example is *meta-control logic*, based on inspecting the (results of) multiple DCPs to take decisions about their evolution.

4.4. Implementation in FCPP. We now briefly describe the implementation of $spawn_{XC}$ in the FCPP C++ framework [Aud20], since it is used in the code examples in the next sections. As typical of the style adopted in FCPP, $spawn_{XC}$ is defined as a number of overloads of a template function. Here we focus on the overload that allows the process function to return a pair containing its output value and the status `nvalue` which specifies for each neighbour δ' (possibly also the device itself) whether the instance should be propagated to δ' or not. The signature of the template function is the following:

```
template<typename node_t, typename G, typename S, typename... Ts,
        typename K = typename std::decay_t<S>::value_type,
        typename T = std::decay_t<std::result_of_t<G(K const&, Ts const&...)>>,
        typename R = std::decay_t<tuple_element_t<0,T>>,
        typename B = std::decay_t<tuple_element_t<1,T>>>
std::enable_if_t< std::is_same< B, nvalue<bool> >::value,
                 std::unordered_map< K, R, common::hash<K> > > >
spawn(node_t& node, trace_t call_point, G&& process, S&& key_set, Ts const&... xs)
```

The first line of the signature specifies the type parameters of the function: `node_t` (device object), `G` (process function), `S` (container of keys of process instances), and the variadic types `Ts` of other arguments to the process function. The subsequent lines declare further types deduced from the parameters: `K` (type of the keys in the container `S`), `T` (full return type of the process function), `R` (output value returned by the process function), and `B` (propagation value returned by the process function).

Then, a further constraint is declared on these types through `std::enable_if_t`. In particular the `is_same` constraint requires that type `B` should be an nvalue of Booleans, i.e., that this particular overload of `spawnXC` can propagate process instances differently to different neighbours.

Despite the complexity of the signature, due to the flexibility offered by templates, a call to the `spawnXC` in FCPP can be as simple as:

```
auto r = spawn(CALL, proc_func, new_instances, xarg1, ..., xargk)
```

where `CALL` is a FCPP macro that automatically provides arguments `node` and `call_point`, `proc_func` is a lambda or pointer to the process function, `new_instances` is a collection of process instance keys to be spawned, and the `xargs` are the optional extra arguments to `proc_func`.

In FCPP, the code for the process function of the multi-gradient DCP presented in Example 4.2 is translated as follows:

```
FUN real_t gradient(ARGS, bool isSource) { CODE
  return exchange(CALL, INF, [&](nvalue<real_t> n){
    return min_hood(CALL, n + node.nbr_dist(), mux(isSource, 0, INF));
  });
}

FUN std::unordered_map<device_t, real_t>
multi_gradient(ARGS, bool isSource, real_t theta) { CODE
  common::option<device_t> gen_set;
  if (isSource) gen_set = node.uid;
  return spawn(CALL, [&](device_t i) {
    real_t output = gradient(CALL, uid() == i);
    nvalue<bool> status = output <= theta && nbr(CALL, node.uid) != i;
    return make_tuple(output, status);
  }, gen_set);
}
```

As apparent, the code has the same structure as the code in Example 4.2, except for the addition of explicit types to variables according to the standard C++ syntax, and the addition of the FCPP-specific macros for passing the context through function calls. Recall that, in C++, syntax `ody[&](args)b` is for anonymous functions.

5. CASE STUDIES

In this section, we present in detail two case studies to illustrate the application of the techniques described in the paper to realistic problems. We cover first the case study on multi-hop message propagation (Section 5.1) and then the case study on spatial monitoring (Section 5.2): for both of them, we discuss the implementation, the experimental setup (scenario, parameters, metrics), and finally the results. Both the use cases have been implemented exploiting the FCPP simulator [Aud20, ART22], which has been extended to support the XC and, in particular, the `spawnXC` built-in construct described above. The experimental framework is publicly available¹, with build infrastructure and instructions, for inspection and reproducibility.

¹<https://github.com/fcpp-experiments/xc-processes>

```

message_log_type r = spawn(CALL, [&](message const& m) {
    nvalue<bool> nv_status = false; // nvalue ⊥[]
    bool target_node = (m.to == node.uid);
    int round = counter(CALL); // just counts the number of rounds run by the device

    if (not target_node and round <= 2) {
        nv_status = mod_self(CALL, nv_status, true);
        nv_status = mod_other(CALL, nv_status, round==1);
    }

    return make_tuple(node.current_time(), nv_status);
}, key_set);

```

Figure 8: FCPP code corresponding to the logic of the DCP for the message propagation scenario.

5.1. Message Propagation. We consider a network of devices where, at some point in time, a source device δ_{FROM} sends a message through a DCP to reach a destination device δ_{TO} . The goal is to show that the DCP mechanism *works* for the specific use case, that multiple strategies exhibiting different dynamics are possible, and that the XC implementation can unlock performance benefits with respect to field calculus-based implementations.

5.1.1. Implementation. We have experimented two different ways of propagation of the DCP. The first one is a *spherical* propagation, where the message originating in δ_{FROM} spreads radially in 3D trying to reach δ_{TO} . The process function executed by each process node implements the following logic:

- (1) if the self device δ is δ_{TO} just return a \perp nvalue $F[]$ (no propagation, since destination has been reached);
- (2) if it is the first round that δ executes the process, and the neighbours it knows (i.e., that have propagated the process to δ) are $\delta_1, \dots, \delta_k$, it propagates it to itself and to its *new neighbours* (that are not yet in the process) by returning an nvalue $T[\delta \mapsto T, \delta_1 \mapsto F, \dots]$;
- (3) if it is the second round that δ executes the process, it propagates the process only to itself
- (4) finally, at the third round δ exits itself the process by returning a \perp nvalue $F[]$.

The snippet of FCPP code in Figure 8 shows the call to $spawn_{XC}$ that implements the DCP described above.

The process identifiers of the DCP instances are `message` objects, which specify the source device `from`, the target device `to`, and the payload. Note that, in the body of the process function: flag `target_node` is \top only on device δ_{TO} ; `nv_status` is the nvalue that determines process propagation; and a call to `mod_self` (resp. `mod_other`) sets the value for the current device (resp. the default value, excluding current neighbours) in an nvalue.

The second strategy is a propagation on a path along a *tree*, where the message originating in δ_{FROM} spreads along a spanning tree of the network of devices to reach δ_{TO} . It is worth pointing out that the spanning tree itself is computed in a fully distributed way by the XC program running on each device. Here, we are just interested in the fact that during such computation each node δ holds three nvalues that are used by $spawn_{XC}$ (see the snippet below):

```

message_log_type r = spawn(CALL, [&](message const& m) {
    nvalue<bool> nv_status = false;
    bool target_node = (m.to == node.uid);
    int round = counter(CALL);

    if (not target_node and round <= 2) {
        nvalue<bool> nv_up = map_hood([&](device_t d) {return d == parent;}, nv_id);
        nvalue<bool> nv_down = map_hood([&](device_t d) {return d == node.uid;}, nv_parent)
            and map_hood([&](set_t s) {return (s.count(m.to) > 0);}, nv_below);

        nv_status = nv_up or nv_down;
        nv_status = mod_self(CALL, nv_status, round == 1);
    }

    return make_tuple(node.current_time(), nv_status);
}, key_set);

```

Figure 9: FCPP code for the message propagation along a spanning tree.

- **nv_id**, containing the ids of all the neighbours δ' of δ (including δ itself);
- **nv_parents**, containing the ids of the parents of all the neighbours δ' of δ in the spanning tree;
- **nv_below**, containing, for each neighbour δ' , the set of ids of the nodes below δ' in the spanning tree.

The process function executed by each process node implements the following logic:

- (1) as in the spherical propagation, if the self device δ is δ_{TO} just return a \perp nvalue $F[]$ (no propagation, since destination has been reached);
- (2) if it is the first or second round that δ executes the process, two nvalues are computed: **nv_up** (resp. **nv_down**) where each neighbour is associated with \top iff it is the parent of δ in the path from δ_{FROM} to the root of the spanning tree (resp. it is the child of δ in the path from the root of the spanning tree to δ_{TO}); taking the logical *or* of these nvalues yields an nvalue **nv_status** which associates a \top value only to the *next* device δ' on the path from δ_{FROM} to δ_{TO} along the spanning tree
 - (a) if it is the first round, δ is added to **nv_status** to propagate the DCP just to devices δ and δ' ;
 - (b) if it is the second round, δ is no longer added to **nv_status**, and the DCP only propagates to δ' ;

The code in Figure 9 shows the call to $spawn_{XC}$ that implements the DCP described above. The part that differs from the previous case is the code executed by δ at the first and second round, in case δ is not the target of the message. The two nvalues **nv_up** and **nv_down** explained above are computed by applying the `map_hood` builtin function to the nvalues **nv_id**, **nv_parent**, and **nv_below** resulting from the computation of the spanning tree. As suggested by the name, `map_hood` transforms an nvalue into another nvalue (possibly of different type) by applying a given function to each element.

5.1.2. *Setup: Scenario, Parameters, Metrics.* For simplicity, we assume devices to be stationary in the simulation (*speed* equal to 0), thus inducing a fixed network topology (although the proof-of-concept program could be run with dynamic topologies as well). We consider

a network of 578 devices into a volume of approximately $1350m \times 1350m \times 20m$, with a connection range of $100m$; so that the network spans about 20 *hops* and devices have an average neighbourhood size (*dens*) of 10. Devices periodically compute the program about every second for 50 seconds of simulated time, with a standard deviation on time intervals (*tvar*) of 5%. We measure the following metrics over the time of execution of the simulations:

- (1) the average number of active processes *aprocc* within the network of devices
- (2) the total number of messages *dcount* correctly delivered to their destinations

The metrics shown in the figures below are obtained by averaging on 100 executions.




Figure 10: Average number of active processes over time for the single-process use case and spherical message propagation.




Figure 11: Average number of active processes over time for the multi-process use case and spherical message propagation.

5.1.3. *Results.* We exploit FCPP to simulate a first use case where only one process is generated. The simulation shows that the process propagates *as a wave* starting from δ_{FROM} outwards. Immediately after the wave front goes beyond a device, the device itself exits the process thus releasing potentially precious resources for other computations. Figure 10 (left) shows the *aprocc* metric. For the first 10 sec, the average is 0, since no process has

been created yet. After the process is created by δ_{FROM} , it propagates until it reaches its destination, and then quickly vanishes. Moreover, from Figure 10 (right) showing the $dcount$ metric, we can see that the message is always delivered.

In a second configuration of the use case, we let 10 different devices generate a new process at each round with probability 5%, in the time interval $[1, 25]$. Figure 11 (left) shows that $aprocs$ grows from 0 (at time 0) up to about 1.6 (approximately at times 10 and 25), and then quickly drops again down to 0. More than 12 processes are generated during the use case, see the $dcount$ metric in Figure 11 (right); however, the average number of active processes $aprocs$ is kept low by the fact that the nodes exploit $spawn_{XC}$ to immediately exit processes after entering and propagating them.




Figure 12: Average number of active processes over time and total delivered messages for the tree propagation use case. Comparison between $spawn_{FC}$ and $spawn_{XC}$.

Figure 12 shows the $aprocs$ and $dcount$ metrics for the tree propagation use case, when only one process is created. The red graphics were generated using $spawn_{XC}$. It is clear that the $aprocs$, which tops at about 0.01, is much smaller than in the case of spherical propagation, which tops at more than 0.7 (Figure 10). In a real scenario, a tradeoff between the reliability and efficiency offered by the two strategies should be considered.

We have also computed the metrics using $spawn_{FC}$, i.e., the FC-based spawn operator also available in FCPP (blue dashed graphics). In order to compare as fairly as possible between XC and FC, we have adopted a wave-like termination for the processes generated by $spawn_{FC}$ [ACT24]. Again, we note a more than 10-fold reduced maximum resource-consumption when using $spawn_{XC}$, since $aprocs$ tops at 0.14 for $spawn_{FC}$. Therefore, the ability of XC to return different states for different neighbours pays off in this scenario.

5.2. Monitoring through Replication. The second scenario of our evaluation has the main purpose of illustrating another archetypal usage of distributed collective processes, that of *replicated algorithms* [PBV16], applying it in the context of *distributed monitoring* [ADS⁺22, ACD⁺21b].

$\phi ::= \perp \mid \top \mid q \mid (\neg\phi) \mid (\phi \wedge \phi) \mid (\phi \vee \phi) \mid (\phi \Rightarrow \phi) \mid (\phi \Leftrightarrow \phi)$	logical
$\mid (\mathsf{P}\phi) \mid (\mathsf{AP}\phi) \mid (\mathsf{EP}\phi) \mid (\mathsf{H}\phi) \mid (\mathsf{AH}\phi) \mid (\mathsf{EH}\phi)$	temporal
$\mid (\mathsf{Y}\phi) \mid (\mathsf{AY}\phi) \mid (\mathsf{EY}\phi) \mid (\phi \mathsf{S}\phi) \mid (\phi \mathsf{AS}\phi) \mid (\phi \mathsf{ES}\phi)$	
<hr style="width: 50%; margin: 0 auto;"/>	
$\phi ::= \perp \mid \top \mid q \mid (\neg\phi) \mid (\phi \wedge \phi) \mid (\phi \vee \phi) \mid (\phi \Rightarrow \phi) \mid (\phi \Leftrightarrow \phi)$	logical
$\mid (\Box\phi) \mid (\Diamond\phi) \mid (\partial\phi) \mid (\partial^-\phi) \mid (\partial^+\phi)$	spatial
$\mid (\phi \mathcal{R}\phi) \mid (\phi \mathcal{T}\phi) \mid (\phi \mathcal{U}\phi) \mid (\mathcal{G}\phi) \mid (\mathcal{F}\phi)$	

Figure 13: Syntax of past-CTL (top) and SLCS (bottom).

5.2.1. *Implementation.* In this case study, we consider a network of devices that may experience some critical situation, modelled by a predicate *critic*. In order to monitor the occurrence of such situations, distributed monitors need to be put in place. Previous works introduce two different abstract logics that can conveniently express such monitors, while granting an automatic translation into an aggregate program realising them: Past-CTL [LS95, GMR17, ADS⁺22] and the Spatial Logic of Closure Spaces (SLCS) [CLLM16, CLLM14, ACD⁺21b]. The former allows to express properties about the past, such as $\mathsf{EP}\phi$, to be read as “ ϕ has been true at some point in the past for some device in the network”. However, Past-CTL does not capture a notion of what is happening through space at the *present time*. This is possible with the latter logic SLCS, which comprises predicates such as $\mathcal{F}\phi$, to be read as “ ϕ is currently true in some point in space”.

Figure 13 illustrates the syntax of the past-CTL and SLCS logics (as presented in [ACD⁺21b]), which build on atomic propositions q to represent observables, using standard logical operators alongside additional temporal or spatial modalities.

In this framework, past-CTL formulas ϕ can be interpreted on augmented event structures \mathbb{E} (Figure 1), yielding a truth value in each event ϵ . Temporal modalities mirror traditional CTL but with two key distinctions:

- Temporal operators apply to actual (*message*) *paths* in \mathbb{E} , involving events in the past of ϵ (instead of hypothetical future branches), ensuring formulas have a definite runtime-computable truth value;
- Operators come in both quantified and unquantified forms, with the latter referring to the linear past on the same device as the one of ϵ , behaving similarly to past-LTL operators.

The modalities are inspired by terms like *Yesterday*, *Since*, *Previously*, *Historically*. We choose Y , EY , S , AS , ES as primitive modalities, with the following informal meaning:

- $\mathsf{Y}\phi$ means “ ϕ held in the (immediately) previous event $p(\epsilon)$ on the same device $d(\epsilon)$ as ϵ ”;
- $\mathsf{EY}\phi$ means “ ϕ held in some neighbour event of ϵ ”;
- $\phi \mathsf{S}\psi$ means “ ψ held in some past event on the same device $d(\epsilon)$ as ϵ , and ϕ has held in events on $d(\epsilon)$ since then”;
- similarly, $\phi \mathsf{AS}\psi$ (resp. $\phi \mathsf{ES}\psi$) mean “for all (message) paths (resp. exists a path) in \mathbb{E} reaching the current event, ψ held in some event of the path and ϕ has held on the path events since then”.

The other operators can be derived from them through $\mathsf{AY}\phi \triangleq \neg \mathsf{EY}\neg\phi$; $\mathsf{P}\phi \triangleq \top \mathsf{S}\phi$ (similarly for AP , EP with AS , ES); $\mathsf{H}\phi \triangleq \neg \mathsf{P}\neg\phi$ (similarly for AH , EH with EP , AP).

SLCS, on the other hand, is a modal logic where formulas are interpreted on nodes of a (possibly directed) graph, without a time component, thus predicating properties of space. When applied to a real-time distributed system, we can informally understand SLCS formulas as predicating properties of the *current* state of the network, according to the knowledge that is available on each device (that may not be perfect). This is formalised by the notion of *self-stabilisation* [ACD⁺21b]: monitors of SLCS formulas converge to the formula value in a finite time after the network and atomic propositions stabilise. Spatial modalities are divided into *local* modalities (predicating properties of the neighbourhood of each device) and *global* modalities (predicating properties of the whole network). The local modalities include:

- $\Box\phi$ (interior): true at points where all neighbours satisfy ϕ ;
- $\Diamond\phi$ (closure): true at points where a neighbour satisfies ϕ ;
- ∂ , ∂^- and ∂^+ represent respectively *boundary* (closure without interior), *interior boundary* (set without the interior) and *closure boundary* (closure without the set).

We take \Diamond as the primitive modality and define the others through $\Box\phi \triangleq \neg\Diamond\neg\phi$, $\partial\phi \triangleq (\Diamond\phi) \wedge \neg(\Box\phi)$, $\partial^-\phi \triangleq \phi \wedge \neg(\Box\phi)$, $\partial^+\phi \triangleq (\Diamond\phi) \wedge \neg\phi$. The global modalities are:

- $\phi\mathcal{R}\psi$ (reaches): true at the ending points of (finite) paths in the graph whose starting point satisfies ψ , and where ϕ holds on each point on the path;
- $\phi\mathcal{T}\psi$ (touches): true at the end of paths whose start satisfies ψ and where ϕ holds in the rest of the path;
- $\phi\mathcal{U}\psi$ (surrounded by): true at points in an area satisfying ϕ , whose closure boundary satisfies ψ ;
- $\mathcal{G}\phi, \mathcal{F}\phi$ (everywhere, somewhere): true in points such that ϕ holds in every (resp. some) point of every (resp. some) incoming path (thus, true in points such that ϕ holds everywhere/somewhere in the connected component of the point).

We take \mathcal{R} as primitive, and express the others by means of equivalences $\phi\mathcal{T}\psi \triangleq \phi\mathcal{R}\Diamond\psi$, $\phi\mathcal{U}\psi \triangleq \phi \wedge \Box\neg(\neg\psi\mathcal{R}\neg\phi)$, $\mathcal{F}\phi \triangleq \top\mathcal{R}\phi$, $\mathcal{G}\phi \triangleq \neg\mathcal{F}\neg\phi$.

Since information from space cannot be accessed directly, and is only available by exchanging and propagating data, the two logics described above have very distinct features: while Past-CTL logic formulas are computed *exactly* by their aggregate monitors, SLCS formulas are only computed *approximately*. In particular, in [ACD⁺21b] the whole logic is implemented in XC based on the modality \mathcal{F} (“somewhere”), which is implemented approximately through a distributed distance calculation.

In this scenario, we compared the effectiveness of the predicates EP *critic* (ever critic) and \mathcal{F} *critic* (now critic SLCS) with a new approach in computing whether a Boolean value is somewhere true, by replicating Past-CTL operators. We thus wrote a generic algorithm replicator in FCPP:

```

GEN(F) auto replicate(ARGS, F fun, size_t n, times_t t) { CODE
    size_t now = shared_clock(CALL) / t;
    auto res = spawn(CALL, [&](size_t i){
        return make_tuple(fun(), i > now - n);
    }, common::option<size_t>{now});
    for (auto x : res) if (x.first > now - n) now = min(now, x.first);
    return res[now];
}

```

This function creates a new replica of an algorithm `fun` every `t` seconds, dropping older replicas so that no more than `n` are present at each time. The result of the replication is

the result of the oldest replica of the algorithm still alive. By carefully tuning the t and n parameters, we can ensure that the value returned is that of a replica that had just enough time to reach convergence, giving us the most recent view of what is the value of `fun` on recent values in the network. By replicating the EP operator from Past-CTL, we can thus turn “*at some point in the past*” into another implementation of “*somewhere right now*”:

```
FUN bool somewhere(ARGS, bool f, size_t replicas, real_t diameter, real_t infospeed) { CODE
  return replicate(CALL, [&]() {
    return logic::EP(CALL, f);
  }, replicas, diameter / infospeed / (replicas-1));
}
```

Given an estimate D of the network diameter and of the speed v at which information propagates, parameter t is set to $\frac{D}{v(n-1)}$, which ensures that the oldest replica alive in a moment had just enough time to reach convergence.

Remark 5.1 (SLCS variants). In this paper, we follow SLCS as presented in [ACD⁺21b]. As already pointed out in [ACD⁺21b, Remark 2], that presentation is not exactly identical to the one given in [CLLM16, CLLM14]. This is due to a different target application domain: abstract model checking for [CLLM16, CLLM14], and distributed monitoring for [ACD⁺21b]. The different target implies a different choice of the base model (continuous functions on closure spaces vs finite paths on directed graphs), although the two base models can be proven equivalent. Crucially, it also entails a different choice of primitive global operators (surrounded vs reachability), while tweaking the semantics of reachability in order to consider edges in the opposite direction, so that the properties can be computed through local interactions without a global oracle (that is not available in our model).

5.2.2. *Setup: Scenario, Parameters, Metrics.* We consider a wide range of networks, varying the size in terms of (maximum) number of hops (parameter *hops*) between 6 and 16, and varying the average number of devices per neighbourhood (parameter *dens*) between 8 and 18, always with a connection range of 100m. Devices periodically compute the monitors about every second, with a standard deviation of the round time interval (parameter *tvar*) ranging from 0% to 40%, for 100 seconds of simulated time. Devices move following randomly-generated checkpoints in their deployment area, with a *speed* varying from 0 to 20% of the communication range per round. Parameters *tvar*, *dens*, *hops*, and *speed* have the same meaning as explained in Section 5.1.2.

In a first test, we set *tvar*=10, *dens*=10, *hops*=10, and *speed*=10, and measure the average of the devices with a *true* value for four different predicates: *critic*, *ever critic* (EP *critic*), *now critic SLCS* (\mathcal{F} *critic*), *now critic replicated* (*somewhere(critic)*). The number of replicas is set to 4. We let the *critic* predicate become true in the simulated times between 10 and 30 for a single device.

In four additional tests, we measure the percentage of time during executions (the average over all devices) when the computed values of \mathcal{F} *critic* and *somewhere(critic)* are wrong (with value 1 representing a wrong value, and value 0 a correct one). We let parameters *tvar*, *dens*, *hops*, and *speed* vary over ranges of relevant values.




Figure 14: Average truth value of spatio-temporal logic monitors over time.

5.2.3. *Results.* Figure 14 shows the value of the monitors over time. The data points are averaged over 400 runs on networks with random initial positions of devices. All three monitors are very effective in rapidly propagating the information that a critic event is occurring throughout the whole network. Then, the *ever critic* predicate never switches back to false by design. On the other hand, both approaches to computing *now critic* successfully switch back to false some time after the criticality disappears at second 30. The replicated approach, however, has a much faster reaction time, performing the switch around second 40 while the SLCS formula only does it about 80 seconds later, with a much slower reaction time. As explained above, *ideally*, updating the *now critic replicated* predicate when the *critic* goes back to \perp at second 30 would just take the time to broadcast the new value to all the other nodes through multiple hops (i.e., local communications); in other words, it would take more or less the same amount of time as when the *critic* rises to \top at time 20. In practice, the approximate estimation of diameter D and information speed v , and the limited number of replicas cause the observed delay. Instead, updating the *now critic SLCS* (which uses a distributed distance computation under the hood) suffers from the well-known *rising value* problem (also known as *count-to-infinity*) [ADV17], which causes a significantly higher delay compared to *now critic replicated*. This suggests that implementing the *somewhere* operator of SLCS exploiting DCPs would lead to an improved monitoring of SLCS formulas.

Figure 15 shows the percentage of erroneous values computed by monitors during executions, averaged over the network devices and the whole simulation time. Each data point is further averaged over 10 runs on networks with different random seeds. It is easy to see that, for all four parameters *tvar*, *dens*, *hops*, and *speed*, and across the whole ranges in which we make them vary, the *now critic SLCS* monitor is clearly significantly less accurate than the *now critic replicated* one. This confirms that the results of the first test above hold in a very large number of different contexts.

6. DISCUSSION AND EXAMPLES

In this section, we further discuss the proposed abstraction. In particular, we delineate the characteristics of the proposed programming model for DCPs (Section 6.1), and then provide examples of applications (Section 6.2) showing its versatility.




Figure 15: Average error of spatio-temporal logic monitors, varying *tvar* (top left), *dens* (top right), *hops* (bottom left), and *speed* (bottom right).

6.1. Features of the abstraction and programming model.

6.1.1. *Progressive, live, and self-* construction of ensembles.* The DCPs have a *dynamic* domain, that evolves progressively to include more or less devices. The devices at the border of the DCP can choose to expand it to (a subset of) their neighbours and the neighbours themselves can opt in or out. Moreover, since evaluation of the program is repeated over time, the border is *live*, meaning that membership can be always re-evaluated, in order to consider the up-to-date context. Conversely, members that are no longer interested in participating in the collective task, or that have completed the tasks associated to their role, can leave the process by returning `False[]` in the spawn routine, or even start process termination patterns as those investigated in [ACT22].

6.1.2. *Flexible regulation of collective process shape and state.* The shape and state of a DCP can be regulated flexibly, by leveraging different kinds of mechanisms. Two key mechanisms can be exploited to rule the participation of devices into DCPs: (i) the sharing of PIDs with neighbours (i.e., the control of process propagation “from the inside”, namely from the devices of the *internal border* of the DCP), and (ii) the possibility of a device to *quit* a process instance that has been received, preventing further propagation of it (i.e., the control of process shape “from the outside”, namely from the devices of the *external border* of the DCP).

Moreover, the state and shape of a process can be controlled at a collective level, as a result of a collective consensus or decision-making activity. Also, the leader or owner

of the DCP may centralise some decision-making: e.g., it may gather statistics from its members (using adaptive collection algorithms [ACD⁺21a]), and use locally-computed policies to decide whether to let more members join (sharing the local decision with a resilient broadcast algorithm [VBD⁺19]). Between fully centralised and fully decentralised settings, there are intermediate solutions based e.g. on a partitioning of the DCP into sub-groups using partitioning mechanisms that can be applied at the aggregate programming level [AAC⁺23, CPVN19]. The state can be used, for instance, to denote different *phases* of a collective task [CVRA21], hence it is important that all the members of the DCP eventually become aware of the up-to-date situation. Regarding shape control, further flexibility is provided by XC, thanks to differentiated messages to neighbours: this feature could be used to essentially control the direction of process propagation (e.g., by filtering, random selective choice, or any other ad-hoc mechanisms). In future work, it would be interesting to explore how these mechanisms might support morphogenesis in modular robotics.

6.1.3. *Support for privacy-preserving collective computations.* The possibility in XC to send differentiated messages to neighbours (unlike classical field calculi [VBD⁺19]), especially when supported infrastructurally through point-to-point communication channels, can also promote *privacy* in collective computations. By regulating the spread of processes “from the inside”, devices that do not have to be aware of certain tasks are not exposed to such potentially sensible information. Then, of course, a proper management of privacy (and security) would require addressing the issue throughout the software stack, i.e., also at the level of the middleware or platform sustaining the collective computation [PCC⁺18].

6.2. **Examples.** Given its features, the DCP abstraction could turn useful to program several kinds of higher-level distributed computing abstractions and tasks such as, for instance, the following.

6.2.1. *Modelling of teams or ensembles of agents* [NLPT14]. A DCP can represent, through its very domain, the set of agents that belong to a certain team or ensemble. Consider:

```
set<task_t> tasks = /* arbitrary task allocation */;
spawn(CALL, [&](task_t t){
  // add code solving the task
  // run the process where the task is allocated:
  nvalue<bool> nv_out = map_hood([&](device_t id){
    return t.contains(id);
  }, node.nbr_uid());
  return make_tuple(nullptr, nv_out);
}, tasks);
```

Here, there will be a different process instance for each different task, and all the devices assigned to the same task will form a team. Notice, however, that two devices of the team would be able to communicate within the process only if they are connected by a path of devices also belonging to the same team. The DCP associated to a team can spread around the system to gather and (re-)evaluate a membership condition, to effectively recruit agents into different organisational structures [HL04]. The former mechanism is directly supported by our XC implementation, through the notion of *differentiated messages* to neighbours, enabled by nvalues. Concurrent participation to multiple teams is directly supported by the fact that a single device can participate in an arbitrary number of DCPs. As participation




Figure 16: Graphics of interacting DCPs modelling spatiotemporal tuple operations. Each DCP is denoted as a trapezoid-like shape that springs out at a certain event (a round by a single device). Notation: ϵ_{in} and ϵ_{out} mean that the event generates a process modelling an *in* (retrieval) and *out* (writing) tuple operation, respectively; ϵ_M means that a matching *out* tuple for an *in* operation has been found; ϵ_C means that the *out* tuple process has reached consensus about the *in* process to serve; the \dagger superscript denotes a termination event, starting a process to close an existing process.

to multiple DCPs leads to increased resource requirements (in both computation time and message size), the programmer has to take into account performance issues when designing the generation and propagation logic of concurrent DCPs. However, the fully asynchronous and resilient nature of XC implies that some additional slack can be used on top of resource bounds posed by the architecture, as longer round execution or message exchange time (or even a device crash) can be handled seamlessly by the XC programming model. Last but not least, the activity within a DCP can be used to support the coordination *within* the ensemble it represents, e.g., through gossip or information spreading algorithms, whose scope is limited to the domain of the DCP; therefore, it may be useful also for *privacy-preserving* computations.

6.2.2. *Space-based coordination (e.g., spatiotemporal tuples [CVRA21]).* A DCP could also be attached to a spatial location—to implemented *spatially-attached processes*. Consider an application that needs to collectively monitor some environmental quantities in a set of sensible regions of a smart city. This could be encoded as follows:


```

std::vector<region> sensible_regions = { region(...), region(...), ... };
std::vector<region> matching_regions;

for (region const& r : sensible_regions)
    if (nearby(CALL, r)) matching_regions.push_back(r);

spawn(CALL, /* process logic */, matching_regions);

```

These ideas could be leveraged to support space-based coordination, or to implement coordination models like *spatiotemporal tuples* [CVRA21], whereby tuples and tuple operations can be emitted to reside at or query a particular spatial location. To implement the spatiotemporal tuples model, an DCP instance can be used to represent a single out (writing), rd (reading), and in (retrieval) operation—see Figure 16 for a visual example. A tuple is denoted by its out process: it exists as long as its DCP is alive in some device. Creating tuples that reside at a fixed spatial location/area (e.g., as described by geodetic coordinates) or that remain attached to a particular mobile device is straightforward. In the former case, the DCP membership condition is just that the device’s current location is inside or close by the provided spatial location. In the latter case, the DCP membership condition is just that the device’s current distance to the DCP source device (which may be computed by a simple gradient) is within a certain threshold. We may call these *node-attached processes*: as a node moves, a DCP attached to it can follow through, to support collective contextual services; for instance, a node may recruit other nodes and resources for mobile tasks.

6.2.3. Creation of adaptive system structures to support communication and coordination.

The ability of DCPs to capture both the formation evolution and the collective activity of a group of devices within a pervasive computing system can be leveraged to create resilient structures supporting non-local or system-wide coordination. For instance, this can be used to implement messaging channels in a peer-to-peer network of situated neighbour-interacting devices (e.g., in a smart city) [CVA⁺19]: the channel consists only of the devices between the source and the destination of a message, hence avoiding expensive gossip or flooding processes that would (i) consume resources of possibly all the devices in the system, and (ii) exacerbate privacy and security issues. This idea could be implemented as follows:

```

bool source = // ...
bool destination = // ...
real_t width = // ....
bool in_channel = channel(CALL, source, destination, width);
option<device_t> generation_pids;
if (in_channel) generation_pids = broadcast(CALL, source, node.uid);
spawn(CALL, /* process logic */, generation_pids);

```

As another example, consider the *Self-organising Coordination Regions (SCR)* pattern [CPVN19]: it is a mechanism to control the level of decentralisation in systems to flexibly support situated tasks, based on (i) decentralised leader election [PCV22]; (ii) creation of areas around the leaders to basically partition the overall system into manageable regions [WH03]; and (iii) supporting intra-region and inter-region coordination e.g. by means of *information flows* [WH07]. Note that traditional solutions based on field calculi [CPVN19] do not easily allow for the partitions to *overlap*: this, instead, could be desired for fault-tolerance, flexibility, and improved interaction between adjacent regions, and it turns out to

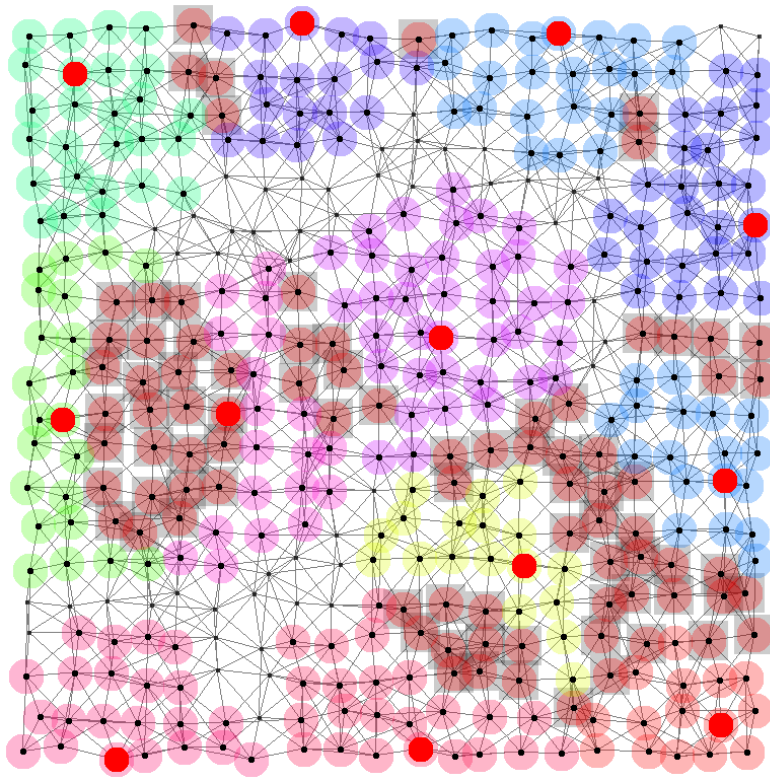


Figure 17: SCR pattern: snapshot from a simulation. Notation: black dots and links represent the network of devices; large red circles denote leaders; semi-transparent coloured circles denote running processes (different colours for different instances) as well as regions; the square semi-transparent shadows are used to highlight devices that run two or more process instances, hence denoting the overlapping of processes.

be easily implementable using DCPs. An implementation of SCR with overlapping regions could be coded as follows:

```

real_t mean_area_size = // ...
bool leader = leader_election(CALL, mean_area_size);
// Now each leader starts a concurrent process that may overlap with others
common::option<device_t> gen;
if (leader) gen = node.uid;
spawn([&](device_t id, bool is_leader, real_t region_size){
    bool region_leader = leader == node.uid and is_leader;
    real_t g = gradient(CALL, region_leader); // cf. Example 3.1
    (region_leader || g < region_size, distributed_sensing(CALL, g))
}, gen, leader, 100.0);

```

In this case, a leader can also act as a regular member of another leaders' region. Figure 17 shows a snapshot from a simulation of this SCR pattern.

7. RELATED WORK

This paper presents the *distributed collective process (DCP)* abstraction and its implementation in XC. In the literature, several macro-level or collective-level abstractions [Cas23b] have been proposed to capture or promote the emergence of coordinated behaviour of whole groups of interacting agents.

Topological and spatial abstractions. These abstractions exploit space as a way to drive collective behaviour. For instance, in *SpatialViews* [NKSI05], groups of nodes (i.e. ensembles) can be matched against spatiotemporal expressions (which define the scope and resolution for selection) and iterated upon to access or command them. In *Abstract Regions* [WM04], regions work as collective communication interfaces. In XC it is also possible to express formulas for matching spatiotemporal regions (e.g., by querying device sensors) and hence the devices therein. In particular, the domain of a DCP fully denotes a spatiotemporal region, with the process logic embedding the rules for setting up, evolving, and tearing down the region; several kinds of shapes and dynamics can be programmed (cf. Section 6). Languages tailored to shape and pattern formation in systems of neighbour-interacting devices have also been proposed, such as *Origami Shape Language (OSL)* [Nag01], where shapes are constructively described as a sequence of straight folds performed on a square sheet; however, these approaches focus on shape creation and are not general-purpose.

Ensembles [NLPT14, BGH⁺13]. Ensembles are dynamic composites of devices. Ensembles can be explicitly or implicitly defined (e.g., dynamically as result of applying attribute-based formation rules), and can be used to structure and refer groups of coordinated entities (e.g., in terms of roles) or to structure coordination itself, for instance by denoting them as targets of communication. In *Distributed Emergent Ensembles of Components (DEECo)* [BGH⁺13], ensembles are formed according to *membership conditions*. In *Service Component Ensemble Language (SCEL)* [NLPT14] and *Attribute-based Communication (AbC)* [AADNL20], ensembles are specified according to predicates on the attributes that the components expose. On the other hand, in Bucchiarone et al. [BM19], ensembles are groups with roles to be filled, and ensemble configurations are described as hypergraphs that specify cells (i.e., roles to be taken by participants), adaptors (i.e., coordination rules), and fragments (provided functionality). There are also practical languages for swarm robotics where ensemble-like abstractions are exposed. In Buzz [PB16], swarms are first-class abstractions that can be handled by union, intersection, iteration, etc. In Voltron [MMWG14], a language for drone sensor networks, there is a notion of a *team* abstraction that can be the target of action commands and that automatically handles the underlying drones according to the (spatiotemporal) tasks at hand. Most of similar approaches focus on task orchestration, leveraging centralised coordinators. For instance, Dolphin [LMPS18] has a notion of vehicle set which is similar to a swarm in Buzz, whereas TeCoLa [KL16], a framework for coordination of robot teams, provides a notion of a “mission group” that is automatically managed according to the sets of services supported by the robots.

Collective communication interfaces. Collective communication interfaces are abstractions for sending data to or receiving data from groups of devices. These abstractions usually enable to flexibly express the targets of communications actions, e.g., via attributes [NLPT14, AADNL20]. Indeed, this abstraction is very much related to the discussed notion of ensemble, as an ensemble can work as a collective communication endpoint.

Collective-based tasks. Collective-based tasks are abstractions keeping track of the lifecycle and state of tasks assigned to whole collectives. These are generally used in *orchestration* approaches. An example is the *SmartSociety* platform [SSV⁺20], which is aimed at heterogeneous collectives involving both humans and machines. In XC, it is also possible to denote collective tasks as regular values, and use gossip-like library functions to distribute them over arbitrary target domains.

Collective data structures. Collective data structures associate data to the members of a collective. A major example is the notion of a *computational field* [MZ04, VBD⁺19]: statically, it is a map from devices to values; by a dynamic perspective, it is a map from events (i.e., devices across space-time) to values. While the long-standing related notion of *artificial potential fields* [War89] associates values with the environment, computational fields associates values with devices. However, it can be noted that, since devices may cover or fill an environment, these two notions may overlap, even though with computational fields there is the idea that devices collaborate to provide local values for these fields to emerge. Computational fields can be used to represent collective inputs/outputs, as well as intermediate values for larger computations; this insight leads to the following notion.

Aggregate computations. Collective computations are the computations performed by collectives of computational entities. Often, these involve interaction among such entities. One major example of collective computations are *aggregate computations* [VBD⁺19], which can be implemented as *field computations* in *field calculi* [VBD⁺19]; in this case, programs are represented as functions mapping input computational fields to output computational fields, while implicitly handling coordination. XC provides a generalisation of field calculi in terms of a single construct, **exchange**, and `nvalues`. The main benefit of these approaches lies in the possibility of representing whole collective behaviours as reusable functions and using functional composition to effectively compose larger collective behaviours from smaller ones.

Aggregate processes. Aggregate processes are dynamic *aggregate computations* [VBD⁺19] on evolving domains of devices, formalised within the FC [ACT21, ACT22, CVA⁺19, CVA⁺21]. The DCP abstraction presented in this work (cf. Section 4.3) is essentially a generalisation of aggregate processes for the XC. The implementation in XC is also more expressive, for it enables to *control how processes spread from a node to neighbours*, rather than controlling shape evolution based on *opt-out* (i.e., the receiving device decides to quit the process) which is the only option in aggregate processes.

8. CONCLUSION AND FUTURE WORK

In this article, we have described the *distributed collective process (DCP)* abstraction, which aims to structure dynamic and concurrent collective activities on ensembles of neighbour-interacting devices running asynchronously in sense–compute–interact rounds. In particular, DCPs model decentralised collective tasks that can spread, shrink, and move across the devices of the collective system that sustain their execution. We have presented the abstraction in the general framework of event structures and provided a formalisation in the *eXchange Calculus (XC)*. XC is a minimal, functional core language for self-organising systems based on *Neighbouring Values*, which supports fine-grained management of DCPs by precisely controlling information sharing within a device’s neighbourhood. The implementation consists of a construct `spawnXC` that handles the generation and execution of DCP instances.

The construct has been included in the FCPP language and the mechanism has been studied using two case studies of message propagation and distributed monitoring, showing its functionality and versatility. A discussion and examples have also been provided to convey how DCPs can be programmed to support various kinds of self-organisation patterns.

The proposed abstraction and implementation can be further developed along different directions. On one hand, DCPs could be more deeply analysed by a dynamical perspective, e.g. by studying the conditions for expansion and termination across different instantiations of the execution model. Then, its suitability for peculiar distributed computing scenarios such as epidemic modelling and morphogenesis could be investigated. On a more practical side, common patterns of DCPs usage could be encapsulated into a library of reusable functions, to foster application development.

ACKNOWLEDGMENT

This publication is part of project NODES, which has received funding from the MUR – M4C2 1.5 of PNRR with grant agreement no. ECS00000036. The work was also supported by the Italian PRIN project “CommonWears” (2020HCWWLP).

REFERENCES

- [AAC⁺23] Gianluca Aguzzi, Giorgio Audrito, Roberto Casadei, Ferruccio Damiani, Gianluca Torta, and Mirko Viroli. A field-based computing approach to sensing-driven clustering in robot swarms. *Swarm Intell.*, 17(1):27–62, 2023. doi:10.1007/s11721-022-00215-y.
- [AADNL20] Yehia Abd Alrahman, Rocco De Nicola, and Michele Loreti. Programming interactions in collective adaptive systems by relying on attribute-based communication. *Science of Computer Programming*, 192, 2020. doi:10.1016/j.scico.2020.102428.
- [ABDV18] Giorgio Audrito, Jacob Beal, Ferruccio Damiani, and Mirko Viroli. Space-time universality of field calculus. In Giovanna Di Marzo Serugendo and Michele Loreti, editors, *Coordination Models and Languages - 20th IFIP WG 6.1 International Conference, COORDINATION 2018, Held as Part of the 13th International Federated Conference on Distributed Computing Techniques, DisCoTec 2018, Madrid, Spain, June 18-21, 2018. Proceedings*, volume 10852 of *Lecture Notes in Computer Science*, pages 1–20. Springer, 2018. doi:10.1007/978-3-319-92408-3_1.
- [ACD⁺21a] Giorgio Audrito, Roberto Casadei, Ferruccio Damiani, Danilo Pianini, and Mirko Viroli. Optimal resilient distributed data collection in mobile edge environments. *Comput. Electr. Eng.*, 96(Part):107580, 2021. doi:10.1016/j.compeleceng.2021.107580.
- [ACD⁺21b] Giorgio Audrito, Roberto Casadei, Ferruccio Damiani, Volker Stolz, and Mirko Viroli. Adaptive distributed monitors of spatial properties for cyber-physical systems. *J. Syst. Softw.*, 175:110908, 2021. doi:10.1016/J.JSS.2021.110908.
- [ACD⁺22] Giorgio Audrito, Roberto Casadei, Ferruccio Damiani, Guido Salvaneschi, and Mirko Viroli. Functional programming for distributed systems with XC. In Karim Ali and Jan Vitek, editors, *36th European Conference on Object-Oriented Programming, ECOOP 2022, June 6-10, 2022, Berlin, Germany*, volume 222 of *LIPICs*, pages 20:1–20:28. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022. doi:10.4230/LIPICs.ECOOP.2022.20.
- [ACD⁺23] Giorgio Audrito, Roberto Casadei, Ferruccio Damiani, Gianluca Torta, and Mirko Viroli. Programming distributed collective processes for dynamic ensembles and collective tasks. In *Coordination Models and Languages - 25th IFIP WG 6.1 International Conference, COORDINATION 2023, Held as Part of the 18th International Federated Conference on Distributed Computing Techniques, DisCoTec 2023, Lisbon, Portugal, June 19-23, 2023, Proceedings*, volume 13908 of *Lecture Notes in Computer Science*, pages 71–89. Springer, 2023. doi:10.1007/978-3-031-35361-1_4.
- [ACDV23] Giorgio Audrito, Roberto Casadei, Ferruccio Damiani, and Mirko Viroli. Computation against a neighbour: Addressing large-scale distribution and adaptivity with functional programming and scala. *Log. Methods Comput. Sci.*, 19(1), 2023. doi:10.46298/LMCS-19(1:6)2023.

- [ACL⁺20] Moreno Ambrosin, Mauro Conti, Riccardo Lazzeretti, Md Masoom Rabbani, and Silvio Ranise. Collective remote attestation at the internet of things scale: State-of-the-art and future challenges. *IEEE Commun. Surv. Tutorials*, 22(4):2447–2461, 2020. doi:10.1109/COMST.2020.3008879.
- [ACT21] Giorgio Audrito, Roberto Casadei, and Gianluca Torta. Towards integration of multi-agent planning with self-organising collective processes. In *IEEE International Conference on Autonomous Computing and Self-Organizing Systems, ACSOS 2021, Companion Volume, Washington, DC, USA, September 27 - Oct. 1, 2021*, pages 297–298. IEEE, 2021. doi:10.1109/ACSOS-C52956.2021.00042.
- [ACT22] Giorgio Audrito, Roberto Casadei, and Gianluca Torta. On the dynamic evolution of distributed computational aggregates. In *IEEE International Conference on Autonomous Computing and Self-Organizing Systems Companion, ACSOS-C 2022, Virtual, CA, USA, September 19-23, 2022*, pages 37–42. IEEE, 2022. doi:10.1109/ACSOS-C56246.2022.00024.
- [ACT24] Giorgio Audrito, Roberto Casadei, and Gianluca Torta. A general framework and decentralised algorithms for collective computational processes. *Future Gener. Comput. Syst.*, 158:11–27, 2024. doi:10.1016/j.future.2024.04.020.
- [ADS⁺22] Giorgio Audrito, Ferruccio Damiani, Volker Stolz, Gianluca Torta, and Mirko Viroli. Distributed runtime verification by past-CTL and the field calculus. *J. Syst. Softw.*, 187:111251, 2022. doi:10.1016/J.JSS.2022.111251.
- [ADV17] Giorgio Audrito, Ferruccio Damiani, and Mirko Viroli. Optimally-self-healing distributed gradient structures through bounded information speed. In Jean-Marie Jacquet and Mieke Massink, editors, *Coordination Models and Languages - 19th IFIP WG 6.1 International Conference, COORDINATION 2017, Held as Part of the 12th International Federated Conference on Distributed Computing Techniques, DisCoTec 2017, Neuchâtel, Switzerland, June 19-22, 2017, Proceedings*, volume 10319 of *Lecture Notes in Computer Science*, pages 59–77. Springer, 2017. doi:10.1007/978-3-319-59746-1_4.
- [AIHT17] Dario Albani, Joris IJsselmuiden, Ramon Haken, and Vito Trianni. Monitoring and mapping with robot swarms for agricultural applications. In *14th IEEE International Conference on Advanced Video and Signal Based Surveillance, AVSS 2017, Lecce, Italy, August 29 - September 1, 2017*, pages 1–6. IEEE Computer Society, 2017. doi:10.1109/AVSS.2017.8078478.
- [AIM10] Luigi Atzori, Antonio Iera, and Giacomo Morabito. The internet of things: A survey. *Comput. Networks*, 54(15):2787–2805, 2010. doi:10.1016/j.comnet.2010.05.010.
- [AJAM20] Ross Arnold, Jonathan Jablonski, Benjamin Abruzzo, and Elizabeth Mezzacappa. Heterogeneous UAV multi-role swarming behaviors for search and rescue. In *IEEE Conference on Cognitive and Computational Aspects of Situation Management, CogSIMA 2020, Victoria, BC, Canada, August 24-29, 2020*, pages 122–128. IEEE, 2020. doi:10.1109/CogSIMA49017.2020.9215994.
- [ART22] Giorgio Audrito, Luigi Rapetta, and Gianluca Torta. Extensible 3d simulation of aggregated systems with FCPP. In *Coordination Models and Languages - 24th IFIP WG 6.1 International Conference, COORDINATION 2022, Held as Part of the 17th International Federated Conference on Distributed Computing Techniques, DisCoTec 2022, Lucca, Italy, June 13-17, 2022, Proceedings*, volume 13271 of *Lecture Notes in Computer Science*, pages 55–71. Springer, 2022. doi:10.1007/978-3-031-08143-9_4.
- [Aud20] Giorgio Audrito. FCPP: an efficient and extensible field calculus framework. In *IEEE International Conference on Autonomous Computing and Self-Organizing Systems, ACSOS 2020, Washington, DC, USA, August 17-21, 2020*, pages 153–159. IEEE, 2020. doi:10.1109/ACSOS49614.2020.00037.
- [BBHR20] Olivier Boissier, Rafael H Bordini, Jomi Hubner, and Alessandro Ricci. *Multi-agent oriented programming: programming multi-agent systems using JaCaMo*. Mit Press, 2020.
- [BDU⁺13] Jacob Beal, Stefan Dulman, Kyle Usbeck, Mirko Viroli, and Nikolaus Correll. Organizing the aggregate: Languages for spatial computing. In *Formal and Practical Aspects of Domain-Specific Languages: Recent Developments*, chapter 16, pages 436–501. IGI Global, 2013. doi:10.4018/978-1-4666-2092-6.ch016.
- [BFBD13] Manuele Brambilla, Eliseo Ferrante, Mauro Birattari, and Marco Dorigo. Swarm robotics: a review from the swarm engineering perspective. *Swarm Intell.*, 7(1):1–41, 2013. doi:10.1007/s11721-012-0075-2.

- [BGH⁺13] Tomás Bures, Ilias Gerostathopoulos, Petr Hnetynka, Jaroslav Keznikl, Michal Kit, and Frantisek Plasil. DEECO: an ensemble-based component system. In *CBSE'13, Proceedings of the 16th ACM SIGSOFT Symposium on Component Based Software Engineering, part of CompArch '13*, pages 81–90. ACM, 2013. doi:10.1145/2465449.2465462.
- [BM19] Antonio Bucchiarone and Marina Mongiello. Ten years of self-adaptive systems: From dynamic ensembles to collective adaptive systems. In *From Software Engineering to Formal Methods and Tools, and Back - Essays Dedicated to Stefania Gnesi on the Occasion of Her 65th Birthday*, volume 11865 of *Lecture Notes in Computer Science*, pages 19–39. Springer, 2019. doi:10.1007/978-3-030-30985-5_3.
- [Bul14] Nils Bulling. A survey of multi-agent decision making. *KI - Künstliche Intelligenz*, 28(3):147–158, July 2014. doi:10.1007/s13218-014-0314-3.
- [Cas23a] Roberto Casadei. Artificial collective intelligence engineering: A survey of concepts and perspectives. *Artif. Life*, 29(4):433–467, 2023. doi:10.1162/ARTL_A_00408.
- [Cas23b] Roberto Casadei. Macroprogramming: Concepts, state of the art, and opportunities of macroscopic behaviour modelling. *ACM Computing Surveys*, January 2023. doi:10.1145/3579353.
- [CLLM14] Vincenzo Ciancia, Diego Latella, Michele Loreti, and Mieke Massink. Specifying and verifying properties of space. In Josep Díaz, Ivan Lanese, and Davide Sangiorgi, editors, *Theoretical Computer Science - 8th IFIP TC 1/WG 2.2 International Conference, TCS 2014, Rome, Italy, September 1-3, 2014. Proceedings*, volume 8705 of *Lecture Notes in Computer Science*, pages 222–235. Springer, 2014. doi:10.1007/978-3-662-44602-7_18.
- [CLLM16] Vincenzo Ciancia, Diego Latella, Michele Loreti, and Mieke Massink. Model checking spatial logics for closure spaces. *Log. Methods Comput. Sci.*, 12(4), 2016. doi:10.2168/LMCS-12(4:2)2016.
- [CPVN19] Roberto Casadei, Danilo Pianini, Mirko Viroli, and Antonio Natali. Self-organising coordination regions: A pattern for edge computing. In Hanne Riis Nielson and Emilio Tuosto, editors, *Coordination Models and Languages - 21st IFIP WG 6.1 International Conference, COORDINATION 2019, Held as Part of the 14th International Federated Conference on Distributed Computing Techniques, DisCoTec 2019, Kongens Lyngby, Denmark, June 17-21, 2019, Proceedings*, volume 11533 of *Lecture Notes in Computer Science*, pages 182–199. Springer, 2019. doi:10.1007/978-3-030-22397-7_11.
- [CVA⁺19] Roberto Casadei, Mirko Viroli, Giorgio Audrito, Danilo Pianini, and Ferruccio Damiani. Aggregate processes in field calculus. In *Coordination Models and Languages - 21st IFIP WG 6.1 International Conference, COORDINATION 2019, Held as Part of the 14th International Federated Conference on Distributed Computing Techniques, DisCoTec 2019 Proceedings*, volume 11533 of *Lecture Notes in Computer Science*, pages 200–217. Springer, 2019. doi:10.1007/978-3-030-22397-7_12.
- [CVA⁺21] Roberto Casadei, Mirko Viroli, Giorgio Audrito, Danilo Pianini, and Ferruccio Damiani. Engineering collective intelligence at the edge with aggregate processes. *Engineering Applications of Artificial Intelligence*, 97:104081, 2021. doi:10.1016/j.engappai.2020.104081.
- [CVRA21] Roberto Casadei, Mirko Viroli, Alessandro Ricci, and Giorgio Audrito. Tuple-based coordination in large-scale situated systems. In *Coordination Models and Languages - 23rd IFIP WG 6.1 International Conference, COORDINATION 2021, Held as Part of the 16th International Federated Conference on Distributed Computing Techniques, DisCoTec 2021, Valletta, Malta, June 14-18, 2021, Proceedings*, volume 12717 of *Lecture Notes in Computer Science*, pages 149–167. Springer, 2021. doi:10.1007/978-3-030-78142-2_10.
- [DJW20] Rocco De Nicola, Stefan Jähnichen, and Martin Wirsing. Rigorous engineering of collective adaptive systems: special section. *Int. J. Softw. Tools Technol. Transf.*, 22(4):389–397, 2020. doi:10.1007/s10009-020-00565-0.
- [DM82] Luís Damas and Robin Milner. Principal type-schemes for functional programs. In Richard A. DeMillo, editor, *Conference Record of the Ninth Annual ACM Symposium on Principles of Programming Languages, Albuquerque, New Mexico, USA, January 1982*, pages 207–212. ACM Press, 1982. doi:10.1145/582153.582176.
- [GC92] David Gelernter and Nicholas Carriero. Coordination languages and their significance. *Commun. ACM*, 35(2):96–107, 1992. doi:10.1145/129630.376083.

- [GD09] Roderich Groß and Marco Dorigo. Towards group transport by swarms of robots. *Int. J. Bio Inspired Comput.*, 1(1/2):1–13, 2009. doi:10.1504/IJBIC.2009.022770.
- [GMR17] Nicola Gigante, Angelo Montanari, and Mark Reynolds. A one-pass tree-shaped tableau for LTL+past. In Thomas Eiter and David Sands, editors, *LPAR-21, 21st International Conference on Logic for Programming, Artificial Intelligence and Reasoning, Maun, Botswana, May 7-12, 2017*, volume 46 of *EPiC Series in Computing*, pages 456–473. EasyChair, 2017. doi:10.29007/3HB9.
- [Gru08] Tom Gruber. Collective knowledge systems: Where the social web meets the semantic web. *J. Web Semant.*, 6(1):4–13, 2008. doi:10.1016/j.websem.2007.11.011.
- [GRWF16] Hendrik-Jorn Gunther, Raphael Riebl, Lars C. Wolf, and Christian Facchi. Collective perception and decentralized congestion control in vehicular ad-hoc networks. In *2016 IEEE Vehicular Networking Conference, VNC 2016, Columbus, OH, USA, December 8-10, 2016*, pages 1–8. IEEE, 2016. doi:10.1109/VNC.2016.7835931.
- [Gup15] Gopal Gupta. Language-based software engineering. *Sci. Comput. Program.*, 97:37–40, 2015. doi:10.1016/j.scico.2014.02.010.
- [HB10] Jim Hendler and Tim Berners-Lee. From the semantic web to social machines: A research challenge for AI on the world wide web. *Artif. Intell.*, 174(2):156–161, 2010. doi:10.1016/j.artint.2009.11.010.
- [HL04] Bryan Horling and Victor R. Lesser. A survey of multi-agent organizational paradigms. *Knowl. Eng. Rev.*, 19(4):281–316, 2004. doi:10.1017/S0269888905000317.
- [IPW01] Atsushi Igarashi, Benjamin C. Pierce, and Philip Wadler. Featherweight Java: a minimal core calculus for Java and GJ. *ACM Trans. Program. Lang. Syst.*, 23(3):396–450, 2001. doi:10.1145/503502.503505.
- [KL16] Manos Koutsoubelias and Spyros Lalis. TeCoLa: A programming framework for dynamic and heterogeneous robotic teams. In *Proceedings of the 13th International Conference on Mobile and Ubiquitous Systems: Computing, Networking and Services (MobiQuitous 2016)*, pages 115–124. ACM, 2016. doi:10.1145/2994374.2994397.
- [KS21] Vasileios Karagiannis and Stefan Schulte. Distributed algorithms based on proximity for self-organizing fog computing systems. *Pervasive Mob. Comput.*, 71:101316, 2021. doi:10.1016/j.pmcj.2020.101316.
- [Lam78] Leslie Lamport. Time, clocks, and the ordering of events in a distributed system. *Commun. ACM*, 21(7):558–565, 1978. doi:10.1145/359545.359563.
- [LHJ19] Chenguang Liu, Jie Hua, and Christine Julien. SCENTS: collaborative sensing in proximity IoT networks. In *IEEE International Conference on Pervasive Computing and Communications Workshops, PerCom Workshops 2019, Kyoto, Japan, March 11-15, 2019*, pages 189–195. IEEE, 2019. doi:10.1109/PERCOMW.2019.8730863.
- [LLM17] Alberto Lluch-Lafuente, Michele Loreti, and Ugo Montanari. Asynchronous distributed execution of fixpoint-based computational fields. *Log. Methods Comput. Sci.*, 13(1), 2017. doi:10.23638/LMCS-13(1:13)2017.
- [LMPS18] Keila Lima, Eduardo R. B. Marques, José Pinto, and João B. Sousa. Dolphin: A task orchestration language for autonomous vehicle networks. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS 2018, Madrid, Spain, October 1-5, 2018*, pages 603–610. IEEE, 2018. doi:10.1109/IROS.2018.8594059.
- [LS95] François Laroussinie and Philippe Schnoebelen. A hierarchy of temporal logics with past. *Theor. Comput. Sci.*, 148(2):303–324, 1995. doi:10.1016/0304-3975(95)00035-U.
- [MDB19] Yuanqiu Mo, Soura Dasgupta, and Jacob Beal. Robustness of the adaptive Bellman-Ford algorithm: Global stability and ultimate bounds. *IEEE Trans. Autom. Control.*, 64(10):4121–4136, 2019. doi:10.1109/TAC.2019.2904239.
- [MK16] Nitinder Mohan and Jussi Kangasharju. Edge-fog cloud: A distributed cloud for internet of things computations. In *2016 Cloudification of the Internet of Things, CIoT 2016, Paris, France, November 23-25, 2016*, pages 1–6. IEEE, 2016. doi:10.1109/CIOT.2016.7872914.
- [MMWG14] Luca Mottola, Mattia Moretta, Kamin Whitehouse, and Carlo Ghezzi. Team-level programming of drone sensor networks. In *Proceedings of the 12th ACM Conference on Embedded Network Sensor Systems (SenSys’14)*, pages 177–190. ACM, 2014. doi:10.1145/2668332.2668353.

- [MWT⁺19] Kimberly McGuire, Christophe De Wagter, Karl Tuyls, H. J. Kappen, and Guido C. H. E. de Croon. Minimal navigation solution for a swarm of tiny flying robots to explore an unknown environment. *Sci. Robotics*, 4(35), 2019. doi:10.1126/scirobotics.aaw9710.
- [MZ04] Marco Mamei and Franco Zambonelli. Programming pervasive and mobile computing applications with the TOTA middleware. In *Proceedings of the Second IEEE International Conference on Pervasive Computing and Communications (PerCom 2004), 14-17 March 2004, Orlando, FL, USA*, pages 263–276. IEEE Computer Society, 2004. doi:10.1109/PERCOM.2004.1276864.
- [Nag01] Radhika Nagpal. *Programmable self-assembly: constructing global shape using biologically-inspired local interactions and origami mathematics*. PhD thesis, Massachusetts Institute of Technology, Cambridge, MA, USA, 2001. URL: <http://hdl.handle.net/1721.1/86667>.
- [NKSI05] Yang Ni, Ulrich Kremer, Adrian Stere, and Liviu Iftode. Programming ad-hoc networks of mobile and resource-constrained devices. In *Proceedings of the ACM SIGPLAN 2005 Conference on Programming Language Design and Implementation, Chicago, IL, USA, June 12-15, 2005*, pages 249–260. ACM, 2005. doi:10.1145/1065010.1065040.
- [NLPT14] Rocco De Nicola, Michele Loreti, Rosario Pugliese, and Francesco Tiezzi. A formal approach to autonomic systems programming: The SCEL language. *ACM Trans. Auton. Adapt. Syst.*, 9(2):7:1–7:29, 2014. doi:10.1145/2619998.
- [NM13] Iñaki Navarro and Fernando Matia. A survey of collective movement of mobile robots. *International Journal of Advanced Robotic Systems*, 10(1):73, January 2013. doi:10.5772/54600.
- [NPW81] Mogens Nielsen, Gordon D. Plotkin, and Glynn Winskel. Petri nets, event structures and domains, part I. *Theor. Comput. Sci.*, 13:85–108, 1981. doi:10.1016/0304-3975(81)90112-2.
- [NW04] Ryan Newton and Matt Welsh. Region streams: Functional macroprogramming for sensor networks. In *Workshop on Data Management for Sensor Networks*, pages 78–87, 2004. doi:10.1145/1052199.1052213.
- [PA98] George A. Papadopoulos and Farhad Arbab. Coordination models and languages. *Adv. Comput.*, 46:329–400, 1998. doi:10.1016/S0065-2458(08)60208-9.
- [PB16] Carlo Pinciroli and Giovanni Beltrame. Buzz: A programming language for robot swarms. *IEEE Softw.*, 33(4):97–100, 2016. doi:10.1109/MS.2016.95.
- [PBV16] Danilo Pianini, Jacob Beal, and Mirko Viroli. Improving gossip dynamics through overlapping replicates. In *Coordination Models and Languages - 18th IFIP WG 6.1 International Conference, COORDINATION 2016, Held as Part of the 11th International Federated Conference on Distributed Computing Techniques, DisCoTec 2016, Heraklion, Crete, Greece, June 6-9, 2016, Proceedings*, volume 9686 of *Lecture Notes in Computer Science*, pages 192–207. Springer, 2016. doi:10.1007/978-3-319-39519-7_12.
- [PCC⁺18] Danilo Pianini, Giovanni Ciatto, Roberto Casadei, Stefano Mariani, Mirko Viroli, and Andrea Omicini. Transparent protection of aggregate computations from byzantine behaviours via blockchain. In *Proceedings of the 4th EAI International Conference on Smart Objects and Technologies for Social Good, GOODTECHS 2018, Bologna, Italy, November 28-30, 2018*, pages 271–276. ACM, 2018. doi:10.1145/3284869.3284870.
- [PCV⁺21] Danilo Pianini, Roberto Casadei, Mirko Viroli, Stefano Mariani, and Franco Zambonelli. Time-fluid field-based coordination through programmable distributed schedulers. *Log. Methods Comput. Sci.*, 17(4), 2021. doi:10.46298/lmcs-17(4:13)2021.
- [PCV22] Danilo Pianini, Roberto Casadei, and Mirko Viroli. Self-stabilising priority-based multi-leader election and network partitioning. In *IEEE International Conference on Autonomic Computing and Self-Organizing Systems, ACSOS 2022, Virtual, CA, USA, September 19-23, 2022*, pages 81–90. IEEE, 2022. doi:10.1109/ACSOS55765.2022.00026.
- [PCVN21] Danilo Pianini, Roberto Casadei, Mirko Viroli, and Antonio Natali. Partitioned integration and coordination via the self-organising coordination regions pattern. *Future Gener. Comput. Syst.*, 114:44–68, 2021. doi:10.1016/j.future.2020.07.032.
- [SSV⁺20] Ognjen Scekic, Tommaso Schiavinotto, Svetoslav Videnov, Michael Rovatsos, Hong Linh Truong, Daniele Miorandi, and Schahram Dustdar. A programming model for hybrid collaborative adaptive systems. *IEEE Trans. Emerg. Top. Comput.*, 8(1):6–19, 2020. doi:10.1109/TETC.2017.2702578.
- [SYN⁺21] Bharath Sudharsan, Piyush Yadav, Duc-Duy Nguyen, Jefkine Kafunah, and John G. Breslin. Ensemble methods for collective intelligence: Combining ubiquitous ML models in IoT. In *2021*

- IEEE International Conference on Big Data (Big Data), Orlando, FL, USA, December 15-18, 2021*, pages 1960–1963. IEEE, 2021. doi:10.1109/BigData52589.2021.9671901.
- [TADT22] Lorenzo Testa, Giorgio Audrito, Ferruccio Damiani, and Gianluca Torta. Aggregate processes as distributed adaptive services for the industrial internet of things. *Pervasive Mob. Comput.*, 85:101658, 2022. doi:10.1016/j.pmcj.2022.101658.
- [VAB⁺18] Mirko Viroli, Giorgio Audrito, Jacob Beal, Ferruccio Damiani, and Danilo Pianini. Engineering resilient collective adaptive systems by self-stabilisation. *ACM Trans. Model. Comput. Simul.*, 28(2):16:1–16:28, 2018. doi:10.1145/3177774.
- [VBD⁺19] Mirko Viroli, Jacob Beal, Ferruccio Damiani, Giorgio Audrito, Roberto Casadei, and Danilo Pianini. From distributed coordination to field calculus and aggregate computing. *J. Log. Algebraic Methods Program.*, 109, 2019. doi:10.1016/j.jlamp.2019.100486.
- [War89] Charles W. Warren. Global path planning using artificial potential fields. In *Proceedings of the 1989 IEEE International Conference on Robotics and Automation, Scottsdale, Arizona, USA, May 14-19, 1989*, pages 316–321. IEEE Computer Society, 1989. doi:10.1109/ROBOT.1989.100007.
- [WG09] Zena Wood and Antony Galton. A taxonomy of collective phenomena. *Applied Ontology*, 4(3-4):267–292, 2009. doi:10.3233/ao-2009-0071.
- [WH03] Danny Weyns and Tom Holvoet. Regional synchronization for simultaneous actions in situated multi-agent systems. In Vladimír Marík, Jörg P. Müller, and Michal Pechoucek, editors, *Multi-Agent Systems and Applications III, 3rd International Central and Eastern European Conference on Multi-Agent Systems, CEEMAS 2003, Prague, Czech Republic, June 16-18, 2003, Proceedings*, volume 2691 of *Lecture Notes in Computer Science*, pages 497–510. Springer, 2003. doi:10.1007/3-540-45023-8_48.
- [WH07] Tom De Wolf and Tom Holvoet. Designing self-organising emergent systems based on information flows and feedback-loops. In *Proceedings of the First International Conference on Self-Adaptive and Self-Organizing Systems, SASO 2007, Boston, MA, USA, July 9-11, 2007*, pages 295–298. IEEE Computer Society, 2007. doi:10.1109/SASO.2007.16.
- [WM04] Matt Welsh and Geoffrey Mainland. Programming sensor networks using abstract regions. In *1st Symposium on Networked Systems Design and Implementation (NSDI 2004), March 29-31, 2004, San Francisco, California, USA, Proceedings*, pages 29–42. USENIX, 2004. URL: <http://www.usenix.org/events/nsdi04/tech/welsh.html>.