

AperTO - Archivio Istituzionale Open Access dell'Università di Torino

The share operator for field-based coordination

This is the author's manuscript

Original Citation:

Availability:

This version is available <http://hdl.handle.net/2318/1711784> since 2019-09-15T15:59:33Z

Publisher:

Springer Verlag

Published version:

DOI:10.1007/978-3-030-22397-7_4

Terms of use:

Open Access

Anyone can freely access the full text of works made available as "Open Access". Works made available under a Creative Commons license can be used according to the terms and conditions of said license. Use of all other works requires consent of the right holder (author or publisher) if not exempted from copyright protection by the applicable law.

(Article begins on next page)

The share Operator for Field-based Coordination^{*}

Giorgio Audrito¹[0000–0002–2319–0375], Jacob Beal²[0000–0002–1663–5102],
Ferruccio Damiani¹[0000–0001–8109–1706], Danilo Pianini³[0000–0002–8392–5409],
and Mirko Viroli³[0000–0003–2702–5702]

¹ Dipartimento di Informatica, University of Torino, Torino, Italy

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

² Raytheon BBN Technologies, Cambridge (MA), USA

jakebeal@ieee.org

³ ALMA MATER STUDIORUM–Università di Bologna, Italy

{danilo.pianini, mirko.viroli}@unibo.it

Abstract. Recent work in the area of coordination models and collective adaptive systems promotes a view of distributed computations as functions manipulating computational fields (data structures spread over space and evolving over time), and introduces the field calculus as a formal foundation for field computations. With the field calculus, evolution (time) and neighbor interaction (space) are handled by separate functional operators: however, this intrinsically limits the speed of information propagation that can be achieved by their combined use. In this paper, we propose a new field-based coordination operator called **share**, which captures the space-time nature of field computations in a single operator that declaratively achieves: *(i)* observation of neighbors’ values; *(ii)* reduction to a single local value; and *(iii)* update and converse sharing to neighbors of a local variable. In addition to conceptual economy, use of the **share** operator also allows many prior field calculus algorithms to be greatly accelerated, which we validate empirically with simulations of a number of frequently used network propagation and collection algorithms.

Keywords: Aggregate Programming · Computational Field · Information Propagation Speed · Spatial Computing

1 Introduction

The number and density of networking computing devices distributed throughout our environment is continuing to increase rapidly. In order to manage and make effective use of such systems, there is likewise an increasing need for software engineering paradigms that simplify the engineering of resilient distributed

^{*} This work has been partially supported by Ateneo/CSP project “AP: Aggregate Programming” (<http://ap-project.di.unito.it/>). This document does not contain technology or technical data controlled under either U.S. International Traffic in Arms Regulation or U.S. Export Administration Regulations.

systems. Aggregate programming [11,37] is one such promising approach, providing a layered architecture in which programmers can describe computations in terms of resilient operations on “aggregate” data structures with values spread over space and evolving in time.

The foundation of this approach is field computation, formalized by the field calculus [36], a terse mathematical model of distributed computation that simultaneously describes both collective system behavior and the independent, unsynchronized actions of individual devices that will produce that collective behavior [8]. Traditionally, in this approach each construct and reusable component is a pure function from fields to fields—a field is a map from a set of space-time computational events to a set of values—and each primitive construct handles just one key aspect of computation: hence, one construct deals with time (i.e., **rep**, providing field evolution) and one with space (i.e., **nbr**, handling neighbor interaction). However, in recent work on the universality of the field calculus, we have identified that the combination of time evolution and neighbor interaction operators in the original field calculus induces a delay, limiting the speed of information propagation that can be achieved efficiently [2].

In this paper, we address this limitation by extending the field calculus with the **share** construct, combining time evolution and neighbor interaction into a single new atomic coordination operator that simultaneously implements: *(i)* observation of neighbors’ values; *(ii)* reduction to a single local value; and *(iii)* update and converse sharing to neighbors of a local variable.

Following a review of the field calculus and its motivating context in Section 2, we introduce the **share** construct in Section 3, empirically validate the predicted acceleration of speed in frequently used network propagation and collection algorithms in Section 4, and conclude with a summary and discussion of future work in Section 5.

2 Background, Motivation, and Related Work

Programming collective adaptive systems is a challenge that has been recognized and addressed in a wide variety of different contexts. Despite the wide variety of goals and starting points, however, the commonalities in underlying challenges have tended to shape the resulting aggregate programming approaches into several clusters of common approaches, as enumerated in [10]: *(i)* “device-abstraction” methods that abstract and simplify the programming of individual devices and interactions (e.g., TOTA [29], Hood [39], chemical models [38], “paintable computing” [13], Meld [1]) or entirely abstract away the network (e.g., BSP [35], MapReduce [18], Kairos [22]); *(ii)* spatial patterning languages that focus on geometric or topological constructs (e.g., Growing Point Language [16], Origami Shape Language [31], self-healing geometries [15,26], cellular automata patterning [40]); *(iii)* information summarization languages that focus on collection and routing of information (e.g., TinyDB [28], Cougar [41], TinyLime [17], and Regiment [32]); *(iv)* general purpose space-time computing models (e.g., StarLisp [27], MGS [20,21], Proto [9], aggregate programming [11]).

The field calculus [36,8] belongs to the last of these classes, the general purpose models. Like other core calculi, such as λ -calculus [14] or π -calculus [30], the field calculus was designed to provide a minimal, mathematically tractable model of computation—in this case with the goal of unifying across a broad class of aggregate programming approaches and providing a principled basis for integration and composition. Indeed, recent analysis [2] has determined that the current formulation of field calculus is space-time universal, meaning that it is able to capture every possible computation over collections of devices sending messages. Field calculus can thus serve as a unifying abstraction for programming collective adaptive systems, and results regarding field calculus have potential implications for all other works in this field.

That same work establishing universality, however, also identified a key limitation of the current formulation of the field calculus, which we are addressing in this paper. In particular, the operators for time evolution and neighbor interaction in field calculus interact such that for most programs either the message size grows with the distance that information must travel or else information must travel significantly slower than the maximum potential speed. The remainder of this section provides a brief review of these key results from [2]: Section 2.1 introduces the underlying space-time computational model used by the field calculus, Section 2.2 provides a review of the field calculus itself, and Section 2.3 explains and illustrates the problematic interaction between time evolution and neighbor interaction operators that will be addressed by the **share** operator in the next section.

2.1 Space-Time Computation

Field calculus considers a computational model in which a program P is periodically and asynchronously executed by each device δ . When an individual device performs a round of execution, that device follows these steps in order: (i) collects information from sensors, local memory, and the most recent messages from neighbors,⁴ the latter in the form of a *neighboring value* map $\phi : \delta \rightarrow \mathbf{v}$ from neighbors to values, (ii) evaluates program P with the information collected as its input, (iii) stores the results of the computation locally, as well as broadcasting it to neighbors and possibly feeding it to actuators, and (iv) sleeps until it is time for the next round of execution. Note that as execution is asynchronous, devices perform executions independently and without reference to the executions of other devices, except insofar as they use state that has arrived in messages. Messages, in turn, are assumed to be collected by some separate thread, independent of execution rounds.

If we take every such execution as an *event* ϵ , then the collection of such executions across space (i.e., across devices) and time (i.e., over multiple rounds) may be considered as the execution of a single aggregate machine with a topology based on information exchanges \rightsquigarrow . The causal relationship between events may then be formalized as defined in [2]:

⁴ Stale messages may expire after some timeout.

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

Figure 1 shows an example of such an event structure, showing how these relations partition events into “causal past”, “causal future”, and non-ordered “concurrent” subspaces with respect to any given event. Interpreting this in terms of physical devices and message passing, a physical device is instantiated as a chain of events connected by \rightsquigarrow relations (representing evolution of state over time with the device carrying state from one event to the next), and any \rightsquigarrow relation between devices represents information exchange from the tail neighbor to the head neighbor. Notice that this is a very flexible and permissive model: there are no assumptions about synchronization, shared identifiers or clocks, or even regularity of events (though of course these things are not prohibited either).

In principle, an execution at ϵ can depend on information from any event in its past and its results can influence any event in its future. As we will see in Section 2.3, however, this is problematic for the field calculus as it has been previously defined.

Our aggregate constructs manipulate then space-time data values (see Figure 1) that map events to values for each event in an event structure:

Definition 2 (Space-Time Value). Let V be any domain of computational values and \mathbf{E} be a given event structure. A space-time value $\Phi = \langle \mathbf{E}, f \rangle$ is a pair comprising the space and a function $f : E \rightarrow V$ that maps the events E of \mathbf{E} to values.

We can then understand an aggregate computer as a “collective” device manipulating such space-time values, and the field calculus as a definition of operations defined both on individual events and simultaneously on aggregate computers.

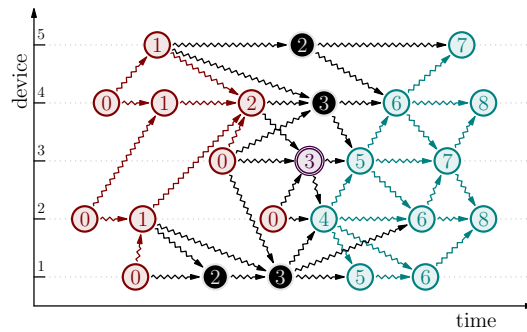


Fig. 1. Example of a space-time event structure, comprising events (circles) and neighbor relations (arrows). Colors indicate causal structure with respect to the doubly-circled event (magenta), splitting events into causal past (red), causal future (cyan) and concurrent (non-ordered, in black). The numbers written within events represent a sample space-time value (cf. Def. 2) associated with that event structure. Figure adapted from [2].

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

Fig. 2. Abstract syntax of the field calculus, adapted from [36]

2.2 Field Calculus

The field calculus is a tiny universal language for computation of space-time values. Figure 2 gives an abstract syntax for field calculus based on the presentation in [36] (covering a subset of the higher-order field calculus in [8], but including all of the issues addressed by the **share** construct). In this syntax, the overbar notation \bar{e} indicates a sequences of elements (e.g., \bar{e} stands for e_1, e_2, \dots, e_n), and multiple overbars are expanded together (e.g., $\bar{\delta} \mapsto \bar{\ell}$ stands for $\delta_1 \mapsto \ell_1, \delta_2 \mapsto \ell_2, \dots, \delta_n \mapsto \ell_n$). There are four keywords in this syntax: **def** and **if** respectively correspond to the standard function definition and the branching expression constructs, while **rep** and **nbr** correspond to the two peculiar field calculus constructs that are the focus of this paper, respectively responsible for evolution of state over time and for sharing information between neighbors.

A field calculus program P is a set of function declarations \bar{F} and the main expression e . This main expression e simultaneously defines both the aggregate computation executed on the overall event structure of an aggregate computer and the local computation executed at each of the individual events therein. An expression e can be:

- A *variable* x , e.g. a function parameter.
- A *value* v , which can be of the following two kinds:
 - a *local value* ℓ , defined via data constructor c and arguments $\bar{\ell}$, such as a Boolean, number, string, pair, tuple, etc;
 - A *neighboring (field) value* ϕ that associates neighbor devices δ to local values ℓ , e.g., a map of neighbors to the distances to those neighbors.
- A function call $f(\bar{e})$ to either a *user-declared function* d (declared with the **def** keyword) or a *built-in function* b , such as a mathematical or logical operator, a data structure operation, or a function returning the value of a sensor.
- A *branching expression* $\text{if}(e_1)\{e_2\}\text{else}\{e_3\}$, used to split a computation into operations on two isolated event structures, where/when e_1 evaluates to **true** or **false**: the result is computation of e_2 in the former area, and e_3 in the latter.

- The `nbr{e}` construct creates a neighboring field value mapping neighbors to their latest available result of evaluating `e`. In particular, each device δ :
 1. shares its value of `e` with its neighbors, and
 2. evaluates the expression into a neighboring field value ϕ mapping each neighbor δ' of δ to the latest value that δ' has shared for `e`.
 Note that within an `if` branch, sharing is restricted to being between device events within the subspace of the branch.
- The `rep(e1){(x) => e2}` construct models state evolution over time: the value of `x` is initialized to `e1`, then evolved at each execution by evaluating `e2`.

Thus, for example, distance to the closest member of a set of “source” devices can be computed with the following simple function:

```
def mux(b, x, y) { if (b) {x} {y} }
def distanceTo(source) {
  rep (infinity) { (d) =>
    mux( source, 0, minHood(nbr{d}+nbrRange()) )
  } }
```

Here, we use the `def` construct to define a `distanceTo` function that takes a Boolean `source` variable as input. The `rep` construct defines a distance estimate `d` that starts at infinity, then decreases in one of two ways. If the `source` variable is true, then the device is currently a source, and its distance to itself is zero. Otherwise, distance is estimated via the triangle inequality, taking the minimum of a neighbor field value (built-in function `minHood`) of the distance to each neighbor (built-in function `nbrRange`) plus that neighbor’s distance estimate `nbr{d}`. Function `mux` ensures that all its arguments are evaluated before being selected.

Additional illustrative examples and full mathematical details of these constructs and the formal semantics of their evaluation can be found in [36].

2.3 Problematic Interaction between rep and nbr Constructs

Unfortunately, the apparently straight-forward combination of state evolution with `nbr` and state sharing with `rep` turns out to contain a hidden delay, which was identified and explained in [2]. This problem may be illustrated by attempting to construct a simple function that spreads information from an event as quickly as possible. Let us say there is a Boolean space-time value `condition`, and we wish to compute a space-time function `ever` that returns true precisely at events where `condition` is true and in the causal future of those events—i.e., spreading out at the maximum theoretical speed throughout the network of devices. One might expect this could be implemented as follows in field calculus:

```
def ever1(condition) {
  rep (false) { (old) => anyHoodPlusSelf(nbr{old}) || condition }
}
```

where `anyHoodPlusSelf` is a built-in function that returns true if any value is true in its neighboring field input (including the value `old` held for the current device). Walking through the evaluation of this function, however, reveals that

there is a hidden delay. In each round, the `old` variable is updated, and will become true if either `condition` is true now for the current device or if `old` was true in the previous round for the current device or for any of its neighbors. Once `old` becomes true, it stays true for the rest of the computation. Notice, however, that a neighboring device does not actually learn that `condition` is true, but that `old` is true. In an event where `condition` first becomes true, the value of `old` that is shared is still false, since the `rep` does not update its value until after the `nbr` has already been evaluated. Only in the next round do neighbors see an updated value of `old`, meaning that `ever1` is not spreading information fast enough to be a correct implementation of `ever`.

We might try to improve this routine by directly sharing the value of `condition`:

```
def ever2(condition) {
  rep (false) { (old) => anyHoodPlusSelf(nbr{old || condition}) }
}
```

This solves the problem for immediate neighbors, but does not solve the problem for neighbors of neighbors, which still have to wait an additional round before `old` is updated.

In fact, it appears that the only way to avoid delays at some depth of neighbor relations is by using unbounded recursion, as previously outlined in [2]:

```
def ever3(condition) {
  rep (false) { (old) =>
    if (countHood() == 0) { old || condition } {
      ever3(anyHoodPlusSelf(nbr{old || condition}))
    } } }
```

where `countHood` counts the number of neighbors, i.e., determining whether any neighbor has reached the same depth of recursion in the branch. Thus, in `ever3`, neighbors' values of `cond` are fed to a nested call to `ever3` (if there are any); and this process is iterated until no more values to be considered are present. This function therefore has a recursion depth equal to the longest sequence of events $\epsilon_0 \rightsquigarrow \dots \rightsquigarrow \epsilon$ ending in the current event ϵ , inducing a linearly increasing computational time and message size and making the routine effectively infeasible for long-running systems.

This case study illustrates the more general problem of delays induced by the interaction of `rep` and `nbr` constructs in field calculus, as identified in [2]. With these constructs, it is never possible to build computations involving long-range communication that are as fast as possible and also lightweight in the amount of communication required.

3 The Share Construct

In order to overcome the problematic interaction between `rep` and `nbr`, we propose a new construct that combines aspects of both:

$$\text{share}(e_1)\{(x) \Rightarrow e_2\}$$

While the syntax of this new **share** construct is identical to that of `rep`, the two constructs differ in the way the construct variable `x` is interpreted each round:

- in **rep**, the value of x is the value produced by evaluating the construct in the previous round, or the result of evaluating e_1 if there is no prior-round value;
- in **share**, on the other hand, x is a *neighboring field* comprising that same value for the current device plus any values of the construct produced by neighbors in their most recent evaluation.

Notice that since x is a neighboring field rather than a local value, e_2 is responsible for processing it into a local value that can be shared with neighbors at the end of the evaluation. Furthermore, notice that the value for δ in the field x corresponds exactly to the value that would be substituted in x for a corresponding **rep** construct. Thus, a **rep** construct may as well be equivalently rewritten as a **share** construct as follows:

$$\text{rep}(e_1)\{(x) \Rightarrow e_2\} \longrightarrow \text{share}(e_1)\{(x) \Rightarrow e_2[x := \text{localHood}(x)]\}$$

where **localHood** is a built-in operator that given a neighboring field ϕ returns the value $\phi(\delta)$ for the current device.

Whenever a field calculus program used x only as **nbr** $\{x\}$ inside the e_2 expression of a **rep**, however, the **share** construct can improve over **rep**. In this case, the following *non-equivalent* rewriting improves the communication speed of an algorithm, while preserving its computational efficiency and overall meaning:

$$\text{rep}(e_1)\{(x) \Rightarrow e_2[\text{nbr}\{x\}]\} \longrightarrow \text{share}(e_1)\{(x) \Rightarrow e_2[x]\}$$

In other words, **share** can be used to *automatically* improve communication speeds of algorithms. Many algorithms with more varied uses of x (e.g., using both x and **nbr** $\{x\}$ in e_2) can be similarly transformed into improved versions.

3.1 Typing and Operational Semantics

Formal typing and operational semantics for the **share** construct is presented in Figure 3 (bottom frame), as an extension to the type system and semantics given in [36, Electronic Appendix]. The typing judgement $\mathcal{A} \vdash e : T$ is to be read “expression e has type T under the set of assumptions \mathcal{A} ”, where \mathcal{A} is a set of assumptions of the form $x : T$ giving type T to variable x . The typing rule [T-SHARE] requires e_1 and e_2 to have the same local (i.e. non-field) type L , assuming x to have the corresponding field type **field**(L), and assigns the same type L to the whole construct.

Example 1 (Typing). Consider the body e of function **ever** as a paradigmatic example (with assumptions $\mathcal{A} = \text{condition} : \text{bool}$):

```
share (false) { (old) => anyHoodPlusSelf(old) || condition }
```

Clearly, $\mathcal{A} \vdash \text{false} : \text{bool}$. Assuming that **anyHoodPlusSelf** is a built-in of type **field**(**bool**) \rightarrow **bool**, we can also conclude that:

$$\mathcal{A}, \text{old} : \text{field}(\text{bool}) \vdash \text{anyHoodPlusSelf}(\text{old}) \mid \text{condition} : \text{bool}.$$

It follows that $\mathcal{A} \vdash e : \text{bool}$.

Value-trees and value-tree environments:	
$\theta ::= \mathbf{v}(\bar{\theta})$	value-tree
$\Theta ::= \bar{\delta} \mapsto \bar{\theta}$	value-tree environment
Auxiliary functions:	
$\phi_0[\phi_1] = \phi_2$ where $\phi_2(\delta) = \begin{cases} \phi_1(\delta) & \text{if } \delta \in \mathbf{dom}(\phi_1) \\ \phi_0(\delta) & \text{otherwise} \end{cases}$	
$\rho(\mathbf{v}(\bar{\theta})) = \mathbf{v}$	
$\pi_i(\mathbf{v}(\theta_1, \dots, \theta_n)) = \theta_i$ if $1 \leq i \leq n$	
$\pi_i(\theta) = \bullet$ otherwise	
For $aux \in \rho, \pi_i$: $\begin{cases} aux(\delta \mapsto \theta) = \delta \mapsto aux(\theta) & \text{if } aux(\theta) \neq \bullet \\ aux(\delta \mapsto \theta) = \bullet & \text{if } aux(\theta) = \bullet \end{cases}$	
$aux(\Theta, \Theta') = aux(\Theta), aux(\Theta')$	
Rules for typing and expression evaluation:	
$\frac{[\mathbf{T-SHARE}] \quad \mathcal{A} \vdash \mathbf{e}_1 : L \quad \mathcal{A}, \mathbf{x} : \mathbf{field}(L) \vdash \mathbf{e}_2 : L}{\mathcal{A} \vdash \mathbf{share}(\mathbf{e}_1)\{\mathbf{x} \Rightarrow \mathbf{e}_2\} : L}$	
$\frac{[\mathbf{E-SHARE}] \quad \begin{array}{l} \delta; \pi_1(\Theta); \sigma \vdash \mathbf{e}_1 \Downarrow \theta_1 \quad \phi' = \rho(\pi_2(\Theta)) \\ \delta; \pi_2(\Theta); \sigma \vdash \mathbf{e}_2[\mathbf{x} := \phi] \Downarrow \theta_2 \quad \phi = (\delta \mapsto \rho(\theta_1))[\phi'] \end{array}}{\delta; \Theta; \sigma \vdash \mathbf{share}(\mathbf{e}_1)\{\mathbf{x} \Rightarrow \mathbf{e}_2\} \Downarrow \rho(\theta_2)\langle \theta_1, \theta_2 \rangle}$	

Fig. 3. Typing and operational semantics for the **share** construct.

The evaluation rule is based on the auxiliary functions given in Figure 3 (middle frame). Function $\rho(\theta)$ extracts the root from a given value-tree, while function $\pi_i(\theta)$ selects the i -th sub-tree of the given value-tree. Both of them can be applied to value-tree environments Θ as well, obtaining a neighboring field (for ρ) or another value-tree environment (for π_i). Furthermore, we use the notation $\phi_0[\phi_1]$ to represent “field update”, so that its result ϕ_2 has $\mathbf{dom}(\phi_2) = \mathbf{dom}(\phi_0) \cup \mathbf{dom}(\phi_1)$ and coincides with ϕ_1 on its domain, or with ϕ_0 otherwise.

The evaluation rule [E-SHARE] produces a value-tree with two branches (for \mathbf{e}_1 and \mathbf{e}_2 respectively). First, it evaluates \mathbf{e}_1 with respect to the corresponding branches of neighbors $\pi_1(\Theta)$ obtaining θ_1 . Then, it collects the results for the construct from neighbors into the neighboring field $\phi' = \rho(\pi_2(\Theta))$. In case ϕ' does not have an entry for δ , $\rho(\theta_1)$ is used obtaining $\phi = (\delta \mapsto \rho(\theta_1))[\phi']$. Finally, ϕ is substituted for \mathbf{x} in the evaluation of \mathbf{e}_2 (with respect to the corresponding branches of neighbors $\pi_2(\Theta)$) obtaining θ_2 , setting $\rho(\theta_2)$ to be the overall value.

Example 2 (Operational Semantics). Consider the body of function **ever**:

```
share (false) { (old) => anyHoodPlusSelf(old) || condition }
```

Suppose that device $\delta = 0$ first executes a round of computation without neighbors (i.e. Θ is empty), and with **condition** equal to **false**. The evaluation of the **share** construct proceeds by evaluating **false** into $\theta_1 = \mathbf{false}()$, gathering neighbor values into $\phi' = \bullet$ (no values are present), and adding the value for the current device obtaining $\phi = (0 \mapsto \mathbf{false})[\bullet] = 0 \mapsto \mathbf{false}$. Finally, the evaluation completes by storing in θ_2 the result of **anyHoodPlusSelf**($0 \mapsto$

`false`)||`false` (which is `false` $\langle \dots \rangle^5$). At the end of the round, device 0 sends a broadcast message containing the result of its overall evaluation, and thus including $\theta^0 = \text{false}\langle \text{false}, \text{false}\langle \dots \rangle \rangle$.

Suppose now that device $\delta = 1$ receives the broadcast message and then executes a round of computation where `condition` is `true`. The evaluation of the `share` constructs starts similarly as before with $\theta_1 = \text{false}\langle \rangle$, $\phi' = 0 \mapsto \text{false}$, $\phi = 0 \mapsto \text{false}$, $1 \mapsto \text{false}$. Then the body of the `share` is evaluated as `anyHoodPlusSelf`($0 \mapsto \text{false}$, $1 \mapsto \text{false}$)||`true` into θ_2 , which is `true` $\langle \dots \rangle$. At the end of the round, device 1 broadcasts the result of its overall evaluation, including $\theta^1 = \text{true}\langle \text{false}, \text{true}\langle \dots \rangle \rangle$.

Then, suppose that device $\delta = 0$ receives the broadcast from device 1 and then performs another round of computation with `condition` equal to `false`. As before, $\theta_1 = \text{false}\langle \rangle$, $\phi = \phi' = 0 \mapsto \text{false}$, $1 \mapsto \text{true}$ and the body is evaluated as `anyHoodPlusSelf`($0 \mapsto \text{false}$, $1 \mapsto \text{true}$)||`false` which produces `true` $\langle \dots \rangle$ for an overall result of $\theta^2 = \text{true}\langle \text{false}, \text{true}\langle \dots \rangle \rangle$.

Finally, suppose that device $\delta = 1$ does not receive that broadcast and discards 0 from its list of neighbor before performing another round of computation with `condition` equal to `false`. Then, $\theta_1 = \text{false}\langle \rangle$, $\phi' = 1 \mapsto \text{true}$, $\phi = (1 \mapsto \text{false})[1 \mapsto \text{true}] = 1 \mapsto \text{true}$, and the body is evaluated as `anyHoodPlusSelf`($1 \mapsto \text{true}$)||`false` which produces `true` $\langle \dots \rangle$.

3.2 The share Construct Improves Communication Speed

To illustrate how `share` solves the problem illustrated in Section 2.3, let us once again consider the `ever` function discussed in that section, for propagating when a `condition` Boolean has ever become true. With the `share` construct, we can finally write a fully functional implementation of `ever` as follows:

```
def ever(condition) {
  share (false) { (old) => anyHoodPlusSelf(old) || condition }
}
```

Function `ever` is simultaneously (i) compact and readable, even more so than `ever1` and `ever2` (note that we no longer need to include the `nbr` construct); (ii) lightweight, as it involves the communication of a single Boolean value each round and few operations; and (iii) optimally efficient in communication speed, since it is true for any event ϵ with a causal predecessor $\epsilon' \leq \epsilon$ where `condition` was true. In particular

- in such an event ϵ' the overall `share` construct is true, since it does so `anyHoodPlusSelf`(old) || `true` regardless of the values in `old`;
- in any subsequent event ϵ'' (i.e. $\epsilon' \rightsquigarrow \epsilon''$) the `share` construct is true since `old` contains a true value (the one coming from ϵ'), and
- the same holds for further following events ϵ by inductive arguments.

⁵ We omit the part of the value tree that are produced by semantic rules not included in this paper, and refer to [36, Electronic Appendix] for the missing parts.

In field calculus alone, such optimal communication speed can be achieved only through unbounded recursion, as argued in [2] and reviewed above in Section 2.3.

The average improvement in communication speed of a routine being converted from the usage of **rep**+**nbr** to **share** according to the rewriting proposed at the beginning of this section can also be statistically estimated, depending on the communication pattern used by the routine.

An algorithm follows a *single-path* communication pattern if its outcome in an event depends essentially on the value of a single selected neighbor: prototypical examples of such algorithms are distance estimations [5,6,4], which are computed out of the value of the single neighbor on the optimal path to the source. In this case, letting T be the average interval between subsequent rounds, the communication delay of an hop is $T/2$ with **share** (since it can randomly vary from 0 to T) and $T/2 + T = 3/2T$ with **rep** + **nbr** (since a full additional round T is wasted in this case). Thus, the usage of **share** allows for an expected three-fold improvement in communication speed for these algorithms.

An algorithm follows a *multi-path* communication pattern if its outcome in an event is obtained from the values of all neighbors: prototypical examples of such algorithms are data collections [3], especially when they are idempotent (e.g. minimums or maximums). In this case, the existence of a single communication path $\epsilon_0 \rightsquigarrow \dots \rightsquigarrow \epsilon$ is sufficient for the value in ϵ_0 to be taken into account in ϵ . Even though the delay of any one of such paths follows the same distribution as for single-path algorithms (0 to T per step with **share**, T to $2T$ per step with **rep** + **nbr**), the overall delay is *minimized* among each existing path. It follows that for sufficiently large numbers of paths, the delay is closer to the minimum of a single hop (0 with **share**, T with **rep** + **nbr**) resulting in an even larger improvement.

4 Application and Empirical Validation

Having developed the **share** construct and shown that it should be able to significantly improve the performance of field calculus programs, we have also applied this development by extending the Protelis [34] implementation of field calculus to support **share** (the implementation is a simple addition of another keyword and accompanying implementation code following the semantics expressed above). We have further upgraded every function in the **protelis-lang** library [19] with an applicable **rep/nbr** combination to use the **share** construct instead, thereby also improving every program that makes use of these libraries of resilient functions. To validate the efficacy of both our analysis and its applied implementation, we empirically validate the improvements in performance for a number of these upgraded functions in simulation.

4.1 Evaluation Setup

We experimentally validate the improvements of the **share** construct through two simulation examples. In both, we deploy a number of mobile devices, computing rounds asynchronously at a frequency of 1 ± 0.1 Hz, and communicating

within a range of 75 meters. All aggregate programs have been written in Protelis [34] and simulations performed in the Alchemist environment [33]. All the results reported in this paper are the average of 200 simulations with different seeds, which lead to different initial device locations, different waypoint generation, and different round frequency. Data generated by the simulator has been processed with Xarray [24] and matplotlib [25]. For the sake of brevity, we do not report the actual code in this paper; however, to guarantee the complete reproducibility of the experiments, the execution of the experiment has been entirely automated, and all the resources have been made publicly available along with instructions.⁶

In the first scenario, we position 2000 mobile devices into a corridor room with sides of, respectively, 200m and 2000m. All but two of the devices are free to move within the corridor randomly, while the remaining two are “sources” that are fixed and located at opposite ends of the corridor. At every point of time, only one of the two sources is active, switching at 80 seconds and 200 seconds (i.e., the active one gets disabled, the disabled one is re-enabled). Devices are programmed to compute a field yielding everywhere the farthest distance from any device to the current active source. In order to do so, they execute the following commonly used coordination algorithms:

1. they compute a potential field measuring the distance from the active source through BIS [6] (`bisGradient` routine in `protelis:coord:spreading`);
2. they accumulate the maximum distance value descending the potential towards the source, through Parametric Weighted Multi-Path C [3] (an optimized version of C in `protelis:coord:accumulation`);
3. they broadcast the information along the potential, from the source to every other device in the system (an optimized version of the `broadcast` algorithm found in `protelis:coord:spreading`, which tags values from the source with a timestamp and propagates them by selecting more recent values).

The choice of the algorithms to be used in validation revealed to be critical. The usage of `share` is able to directly improve the performance of algorithms with solid theoretical guarantees; however, it may also exacerbate errors and instabilities for more ad-hoc algorithms, by allowing them to propagate quicker and more freely, preventing (or slowing down) the stabilization of the algorithm result whenever the network configuration and input is not constant. Of the set of available algorithms for spreading and collecting data, we thus selected variants with smoother recovery from perturbation: optimal single-path distance estimation (BIS gradient [6]), optimal multi-path broadcast [36], and the latest version of data collection (parametric weighted multi-path [3], fine-tuning the weight function).

We are interested in measuring the error of each step (namely, in distance vs. the true values), together with the lag through which these values were generated (namely, by propagating a time-stamp together with values, and computing the

⁶ <https://bitbucket.org/danysk/experiment-2019-coordination-aggregate-share/>

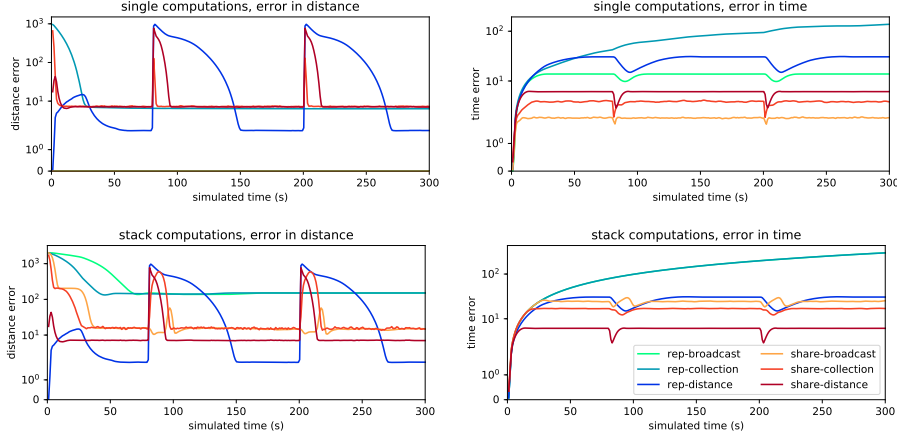


Fig. 4. Performance in the corridor scenario, for both individual algorithms (top) and the composite computation (bottom). Vertical axis is linear in $[0, 1]$ and logarithmic above. Charts on the left column show distance error, while the right column shows time error. The versions of the algorithms implemented with **share** (warm colors) produce significantly less error and converge significantly faster in case of large disruptions than with **rep** (cold colors).

difference with the current time). Moreover, we want to inspect how the improvements introduced by **share** accumulate across the composition of algorithms. To do so, we measure the error in two conditions: (i) composite behavior, in which each step is fed the result computed by the previous step, and (ii) individual behavior, in which each step is fed an ideal result for the previous step, as provided by an oracle.

Figure 4 shows the results from this scenario. Observing the behavior of the individual computations, it is immediately clear how the **share**-based version of the algorithm provides faster recovery from network input discontinuities and lower errors at the limit. These effects are exacerbated when multiple algorithms are composed to build aggregate applications. The only counterexample is the limit of distance estimations, for which **rep** is marginally better, with a relative error less than 1% lower than that of **share**.

Moreover, notice that the collection algorithm with **rep** was not able to recover from changes at all, as shown by the linearly increasing delay in time (and the absence of spikes in distance error). The known weakness of multi-path collection strategies, that is, failing to react to changes due to the creation of information loops, proved to be much more relevant and invalidating with **rep** than with **share**.

In the second example, we deploy 500 devices in a city center, and let them move as though being carried by pedestrians, moving at walking speed ($1.4 \frac{m}{s}$) towards random waypoints along roads open to pedestrian traffic (using map data

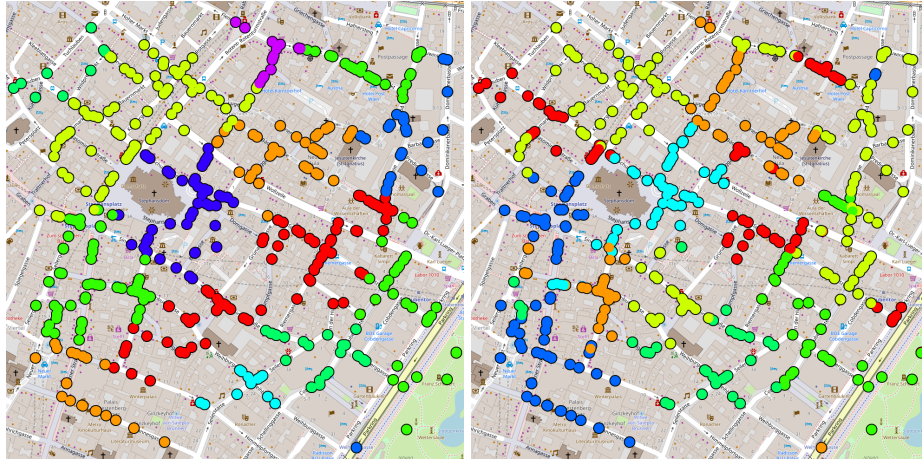


Fig. 5. Snapshots of the Voronoi partitioning scenario using **share** (left) or **rep** (right). Colored dots are simulated devices, with each region having a different color. Faster communication with **share** leads to a higher accuracy in distance estimation, allowing the **share** implementation to perform a better division into regions and preventing regions from expanding beyond their limits.

from OpenStreetMaps [23]). In this scenario, devices must self-organize service management regions with a radius of at most 200 meters, creating a Voronoi partition as shown in Figure 5 (functions `S` and `voronoiPartitioningWithMetric` from `protelis:coord:sparsechoice`). We evaluate performance by measuring the number of partitions generated by the algorithm, and the average and maximum node distance error, where the error for a node n measures how far a node is beyond of the maximum boundary for its cluster. This is computed as $\epsilon_n = \max(0, d(n, l_n) - r)$, where d computes the distance between two devices, l_n is the leader for the cluster n belongs to, and r is the maximum allowed radius of the cluster.

Figure 6 shows the results from this scenario, which also confirm the benefits of faster communication with **share**. The algorithm implemented with **share** has much lower error, mainly due to faster convergence of the distance estimates, and consequent higher accuracy in measuring the distance from the partition leader. Simultaneously, it creates a marginally lower number of partitions, by reducing the amount of occasional single-device regions which arise during convergence and re-organization.

5 Contributions and Future Work

We have introduced a novel **share** construct whose introduction allows a significant acceleration of field calculus programs. We have also made this construct




Fig. 6. Performance in the Voronoi partition scenario: error in distance on the left, leaders count with time on the right. Vertical axis is linear in $[0, 0.1]$ and logarithmic elsewhere. The version implemented with **share** has much lower error: the mean error is negligible, and the most incorrect value, after an initial convergence phase, is close to two orders of magnitude lower than with **rep**, as faster communication leads to more accurate distance estimates. The leader count shows that the systems create a comparable number of partitions, with the **share**-based featuring faster convergence.

available for use in applications though an extension of the Protelis field calculus implementation and its accompanying libraries, and have empirically validated the expected improvements in performance through experiments in simulation.

In future work, we plan to study for which algorithms the usage of **share** may lead to increased instability, thus fine-tuning the choice of **rep** and **nbr** over **share** in the Protelis library. Furthermore, we intend to fully analyze the consequences of **share** for improvement of space-time universality [2], self-adaption [12], and variants of the semantics [7] of the field calculus. It also appears likely that the field calculus can be simplified by the elimination of both **rep** and **nbr** by finding a mapping by which **share** can also be used to implement any usage of **nbr**. Finally, we believe that the improvements in performance will also have positive consequences for nearly all current and future applications that are making use of the field calculus and its implementations and derivatives.

Acknowledgements We thank the anonymous COORDINATION referees for their comments and suggestions on improving the presentation.

References

1. Ashley-Rollman, M.P., Goldstein, S.C., Lee, P., Mowry, T.C., Pillai, P.: Meld: A declarative approach to programming ensembles. In: IEEE International Conference on Intelligent Robots and Systems (IROS '07). pp. 2794–2800 (2007)
2. Audrito, G., Beal, J., Damiani, F., Viroli, M.: Space-time universality of field calculus. In: Coordination Models and Languages. Lecture Notes in Computer Science, vol. 10852, pp. 1–20. Springer (2018)
3. Audrito, G., Bergamini, S., Damiani, F., Viroli, M.: Effective collective summarisation of distributed data in mobile multi-agent systems. In: International Conference on Autonomous Agents and Multiagent Systems (AAMAS). ACM (2019)

4. Audrito, G., Casadei, R., Damiani, F., Viroli, M.: Compositional blocks for optimal self-healing gradients. In: 11th International Conference on Self-Adaptive and Self-Organizing Systems (SASO 2017). pp. 91–100. IEEE (2017)
5. Audrito, G., Damiani, F., Viroli, M.: Optimally-self-healing distributed gradient structures through bounded information speed. In: Coordination Models and Languages. LNCS, vol. 10319, pp. 59–77. Springer (2017)
6. Audrito, G., Damiani, F., Viroli, M.: Optimal single-path information propagation in gradient-based algorithms. *Science of Computer Programming* **166**, 146–166 (2018)
7. Audrito, G., Damiani, F., Viroli, M., Casadei, R.: Run-time management of computation domains in field calculus. In: 1st International Workshops on Foundations and Applications of Self* Systems (FAS*W). pp. 192–197. IEEE (2016)
8. Audrito, G., Viroli, M., Damiani, F., Pianini, D., Beal, J.: A higher-order calculus of computational fields. *ACM Transactions on Computational Logic (TOCL)* **20**(1), 5:1–5:55 (2019)
9. Beal, J., Bachrach, J.: Infrastructure for engineered emergence in sensor/actuator networks. *IEEE Intelligent Systems* **21**, 10–19 (March/April 2006)
10. Beal, J., Dulman, S., Usbeck, K., Viroli, M., Correll, N.: Organizing the aggregate: Languages for spatial computing. In: Formal and Practical Aspects of Domain-Specific Languages: Recent Developments, chap. 16, pp. 436–501. IGI Global (2013)
11. Beal, J., Pianini, D., Viroli, M.: Aggregate programming for the Internet of Things. *IEEE Computer* **48**(9) (2015)
12. Beal, J., Viroli, M., Pianini, D., Damiani, F.: Self-adaptation to device distribution in the Internet of Things. *ACM Transactions on Autonomous and Adaptive Systems (TAAS)* **12**(3), 12:1–12:29 (2017)
13. Butera, W.: Programming a Paintable Computer. Ph.D. thesis, MIT, Cambridge, USA (2002)
14. Church, A.: A set of postulates for the foundation of logic. *Annals of Mathematics* **33**(2), 346–366 (1932)
15. Clement, L., Nagpal, R.: Self-assembly and self-repairing topologies. In: Workshop on Adaptability in Multi-Agent Systems, RoboCup Australian Open (2003)
16. Coore, D.: Botanical Computing: A Developmental Approach to Generating Interconnect Topologies on an Amorphous Computer. Ph.D. thesis, MIT, Cambridge, MA, USA (1999)
17. Curino, C., Giani, M., Giorgetta, M., Giusti, A., Murphy, A.L., Picco, G.P.: Mobile data collection in sensor networks: The tinytime middleware. *Elsevier Pervasive and Mobile Computing Journal* **4**, 446–469 (2005)
18. Dean, J., Ghemawat, S.: Mapreduce: simplified data processing on large clusters. *Communications of the ACM* **51**(1), 107–113 (2008)
19. Francia, M., Pianini, D., Beal, J., Viroli, M.: Towards a foundational api for resilient distributed systems design. In: 2017 IEEE 2nd International Workshops on Foundations and Applications of Self* Systems (FAS* W). pp. 27–32. IEEE (2017)
20. Giavitto, J.L., Godin, C., Michel, O., Prusinkiewicz, P.: Computational models for integrative and developmental biology. Tech. Rep. 72-2002, U. d’Evry, LaMI (2002)
21. Giavitto, J.L., Michel, O., Cohen, J., Spicher, A.: Computations in space and space in computations. In: Unconventional Programming Paradigms, Lecture Notes in Computer Science, vol. 3566, pp. 137–152. Springer, Berlin (2005)
22. Gummadi, R., Gnawali, O., Govindan, R.: Macro-programming wireless sensor networks using kairos. In: Distributed Computing in Sensor Systems (DCOSS). pp. 126–140 (2005)

23. Haklay, M., Weber, P.: OpenStreetMap: User-generated street maps. *IEEE Pervasive Computing* **7**(4), 12–18 (oct 2008)
24. Hoyer, S., Hamman, J.: xarray: N-D labeled arrays and datasets in Python. *Journal of Open Research Software* **5**(1) (2017)
25. Hunter, J.D.: Matplotlib: A 2d graphics environment. *Computing In Science & Engineering* **9**(3), 90–95 (2007)
26. Kondacs, A.: Biologically-inspired self-assembly of 2d shapes, using global-to-local compilation. In: *International Joint Conference on Artificial Intelligence (IJCAI)*. pp. 633–638. Morgan Kaufmann Publishers Inc. (2003)
27. Lasser, C., Massar, J., Miney, J., Dayton, L.: *Starlisp Reference Manual*. Thinking Machines Corporation (1988)
28. Madden, S., Franklin, M.J., Hellerstein, J.M., Hong, W.: TAG: A Tiny AGgregation Service for Ad-hoc Sensor Networks. *SIGOPS Oper. Syst. Rev.* **36**, 131–146 (2002)
29. Mamei, M., Zambonelli, F.: Programming pervasive and mobile computing applications: The tota approach. *ACM Transactions on Software Engineering Methodologies (TOSEM)* **18**(4), 1–56 (2009)
30. Milner, R., Parrow, J., Walker, D.: A calculus of mobile processes, part I. *Information and Computation* **100**(1), 1–40 (September 1992)
31. Nagpal, R.: *Programmable Self-Assembly: Constructing Global Shape using Biologically-inspired Local Interactions and Origami Mathematics*. Ph.D. thesis, MIT, Cambridge, MA, USA (2001)
32. Newton, R., Welsh, M.: Region streams: Functional macroprogramming for sensor networks. In: *Workshop on Data Management for Sensor Networks*. pp. 78–87. DMSN '04, ACM (2004)
33. Pianini, D., Montagna, S., Viroli, M.: Chemical-oriented simulation of computational systems with ALCHEMIST. *J. Simulation* **7**(3), 202–215 (2013)
34. Pianini, D., Viroli, M., Beal, J.: Protelis: Practical aggregate programming. In: *ACM Symposium on Applied Computing 2015*. pp. 1846–1853 (April 2015)
35. Valiant, L.G.: A bridging model for parallel computation. *Communications of the ACM* **33**(8), 103–111 (1990)
36. Viroli, M., Audrito, G., Beal, J., Damiani, F., Pianini, D.: Engineering resilient collective adaptive systems by self-stabilisation. *ACM Transactions on Modelling and Computer Simulation (TOMACS)* **28**(2), 16:1–16:28 (2018)
37. Viroli, M., Beal, J., Damiani, F., Audrito, G., Casadei, R., Pianini, D.: From field-based coordination to aggregate computing. In: *Coordination Models and Languages. Lecture Notes in Computer Science*, vol. 10852, pp. 252–279. Springer (2018)
38. Viroli, M., Pianini, D., Montagna, S., Stevenson, G., Zambonelli, F.: A coordination model of pervasive service ecosystems. *Science of Computer Programming* **110**, 3 – 22 (2015)
39. Whitehouse, K., Sharp, C., Brewer, E., Culler, D.: Hood: a neighborhood abstraction for sensor networks. In: *Proceedings of the 2nd international conference on Mobile systems, applications, and services*. ACM Press (2004)
40. Yamins, D.: *A Theory of Local-to-Global Algorithms for One-Dimensional Spatial Multi-Agent Systems*. Ph.D. thesis, Harvard, Cambridge, MA, USA (2007)
41. Yao, Y., Gehrke, J.: The cougar approach to in-network query processing in sensor networks. *SIGMOD Record* **31**, 9–18 (2002)