

Dolphin: a task orchestration language for autonomous vehicle networks

Keila Lima¹, Eduardo R. B. Marques², José Pinto¹, João B. Sousa¹

Abstract—We present *Dolphin*, an extensible programming language for autonomous vehicle networks. A *Dolphin* program expresses an orchestrated execution of tasks defined compositionally for multiple vehicles. Building upon the base case of elementary one-vehicle tasks, the built-in operators include support for composing tasks in several forms, for instance according to concurrent, sequential, or event-based task flow. The language is implemented as a Groovy DSL, facilitating extension and integration with external software packages, in particular robotic toolkits. The paper describes the *Dolphin* language, its integration with an open-source toolchain for autonomous vehicles, and results from field tests using unmanned underwater vehicles (UUVs) and unmanned aerial vehicles (UAVs).

I. INTRODUCTION

The use of autonomous vehicles is now mainstream for several applications, in particular those making use of several vehicles deployed at once for a common purpose, in networked integration with sensors, human users, and cyberinfrastructures [1–4]. As part of these developments, several software toolkits became popular for networked operation of autonomous vehicles [5–7], allowing for remote control of a single vehicle or basic forms of networked interaction among vehicles. To program a network of autonomous vehicles as an integrated whole, though, we feel that high-level abstractions are required, materialised by the use of domain-specific languages (DSLs) that directly capture the modelling traits of multi-vehicle applications.

In particular, we are concerned with mixed-initiative systems, where humans are part of the control loop and burdened by the intricate complexity of a system-of-systems [8]. In multi-vehicle applications, part of this burden results from the need of separately programming each vehicle without principled mechanisms for coordinated behavior, and the lack of a convenient abstraction for the global state of the system once it is deployed. To attack these problems, we advocate that humans-in-the-loop should be able to write programs that orchestrate the a global behavior of multiple vehicles.

This motivation led us to the development of *Dolphin*, an extensible task orchestration language for autonomous vehicle networks, that is available open-source [9]. A *Dolphin* program expresses an orchestrated execution of tasks defined compositionally for multiple vehicles dynamically available in a network. Building upon the base case of elementary one-vehicle tasks, the built-in operators include support for composing tasks in several forms, for instance

according to concurrent, sequential, or event-based operators, partially inspired by process calculi approaches, e.g., Milner’s CCS [10]. The core language is agnostic and independent of the underlying platform for networked vehicle operations. The system is concretely instantiated through the implementation of abstract programming bindings at the platform level of a robotic toolkit. This is facilitated by the design of *Dolphin* as a Groovy domain-specific language (DSL) [11], allowing direct integration/embedding with/in other Groovy/Java software packages, and seamless addition of extended DSL features.

We developed *Dolphin* bindings for an open-source toolchain [7] developed by Laboratório de Sistemas e Tecnologia Subaquática (LSTS)¹, used to operate heterogeneous types of unmanned vehicles in several experiments over the years (e.g., [12–14]), and bindings for the MAVLink drone protocol [6] are in progress [9]. The LSTS toolchain includes IMC, a message-based interoperability protocol, that has a dedicated subset for the specification and execution of single-vehicle tasks, called IMC plans. Using *Dolphin*, we were able to orchestrate IMC plans in expressive manner for multiple vehicles in field tests using unmanned underwater vehicles (UUVs) and unmanned aerial vehicles (UAVs). We present a field test scenario where three UUVs concurrently perform a bathymetry survey over a given area, and one simulated UAV also in the control loop engaged in a rendezvous maneuver with each of the UUVs at a time.

The rest of the paper is structured as follows. In section II we present *Dolphin* in terms of the underlying architecture, an example scenario, task definition operators, execution engine, and platform bindings. Section III describes the integration of *Dolphin* with the LSTS toolchain and the related tools developed for that purpose. Section IV reports results of field test experiments we conducted for the example scenario using multiple autonomous vehicles. Section V discusses related work. Finally, Section VI ends the paper with concluding remarks and a discussion of future work.

II. THE DOLPHIN LANGUAGE

A. Architecture

The architecture of *Dolphin* is illustrated in Fig. 1. The language engine takes a program supplied by the user and executes it, delegating platform-dependent networked operations to the platform runtime, e.g., polling vehicles in the network or firing tasks for vehicles. Thus, the execution of a *Dolphin* program is centralised, in interface with networked vehicles. We believe this is a convenient

¹ Laboratório de Sistemas e Tecnologia Subaquática, Faculdade de Engenharia da Universidade do Porto, Portugal.

² CRACS/INESC-TEC & Faculdade de Ciências da Universidade do Porto, Portugal.

¹<http://github.com/LSTS>

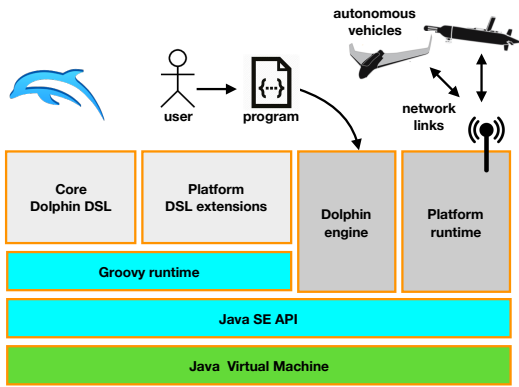


Fig. 1: Dolphin architecture.

approach for mixed-initiative systems, allowing a human to orchestrate an entire network of vehicles based on a global specification. Furthermore, it does not assume the need for peer-to-peer communication among vehicles or any form of tightly coupled interaction among them, even if, of course, these aspects may be relevant for many applications. Later in the paper, we contrast this approach with others (Sec. V) and identify aspects of future work to handle some of the inherent limitations (Sec. VI).

In the Dolphin architecture, the base components support the core DSL embedded in Groovy, and the engine required to execute programs in Java. Platform instantiations may extend the DSL, and must provide a runtime for networked interaction, implementing abstract Java bindings provided by the Dolphin engine. The platform’s DSL extensions and runtime are responsible for implementing suitable constructs for platform tasks and their implementation. The use of Groovy provides a number of features useful for defining the DSL, e.g., operator overloading, meta-class programming, or the use of closures [11]. Moreover, Groovy is fully interoperable with the Java SE API and the Java Virtual Machine.

B. Example program

We now present an example scenario and a corresponding Dolphin program. The scenario at stake, a generalisation of an example given in [15], is illustrated schematically in Fig. 2. It comprises the use of 3 UUVs for joint surveys, executing concurrently over a given area, and of a UAV that responds to the completion of each individual UUV survey by approaching that UUV with a rendezvous behavior (e.g., to retrieve survey data on-the-fly). In Section IV we present an actual configuration and deployment of a variant of this scenario in field tests; here we merely concentrate on its overall meaning and realisation by a Dolphin program.

The Dolphin program is listed in Fig. 3. The code is basically structured in 5 segments: (1) configuration (lines 1–3); (2) vehicle selection (4–14); (3) an auxiliary function for parameterising the UUV tasks (16–24); (4) the execution of desired tasks using all 4 vehicles (25–38), and; (5) program termination, decoupling the vehicles from the program and displaying a simple final message (39–41).

The program begins by asking the user to input a radius




Fig. 2: Example scenario.

```

1 // (1) Configuration
2 r = ask 'Radius of operation area? (km)'
3 APDL = (location 41.18500, -8.70620) ^ r.km
4 // (2) Vehicle selection
5 UUVs = pick {
6     count 3
7     type 'UUV'
8     payload 'DVL','Sidescan'
9     region APDL
10 }
11 UAV = pick {
12     type 'UAV'
13     region APDL
14 }
15 setConnectionTimeout UUVs, 2.min
16 // (3) Function yielding UUV task i
17 def UUVTask ( i ) {
18     imcPlan('survey' + i) >>
19     action { post ready:i } >>
20     imcPlan {
21         planName 'sk' + i
22         keeping duration: 600
23     }
24 }
25 // (4) Execute tasks
26 execute UUVs:
27     UUVTask(1) |
28     UUVTask(2) |
29     UUVTask(3) ,
30     UAV:
31     allOf {
32         when { consume ready:1 }
33         then imcPlan('rv1')
34         when { consume ready:2 }
35         then imcPlan('rv2')
36         when { consume ready:3 }
37         then imcPlan('rv3')
38     }
39 // (5) End
40 release UUVs + UAV
41 message 'Done!'

```

Fig. 3: Dolphin program for the example scenario.

r of a geo-referenced area named **APDL**. The UUVs and UAVs are then selected through two **pick** blocks, one for the UUVs and another one for the UAV. The requirements are that the vehicles are located within the bounds of **APDL**, and, additionally, that each UUV is equipped with specific payload components, a Doppler velocity logger (DVL), and a side-scan sonar. Both the **UUVs** and **UAV** variables stand for vehicle sets (**UAV** is a singleton) that can later be bound to the execution of tasks. Vehicle sets can be manipulated using standard set operators, e.g., $a + b$ (as in line 40 of the program), $a \& b$, and $a - b$ respectively represent the union, intersection, and difference of two sets **a** and **b**.

Following the vehicle selection, a connection timeout of 2 minutes is set for the UUVs using **setConnectionTimeout** (line 15). Up to this time frame, the Dolphin runtime will assume the state of each UUV to be steady in the event that the

connection to it is (intermittently) lost. The timeout setting attends to the fact that UUVs may operate underwater for a long time, during which they may only report state through error-prone and intermittent acoustic communications.

Each UUV task results from the `UUVTask` Groovy function, parameterised by an argument i ; since Dolphin is embedded in Groovy, we may integrate Groovy code within the program at will, making use of standard imperative and object-oriented programming features. The `UUVTask` function yields the sequential composition of three tasks, as specified by the use of the `>>` Dolphin-specific operator: (1) the actual survey task `surveyi`; (2) a notification action signalling readiness for rendezvous, `post ready:i`, that proceeds instantaneously and involves no vehicle interaction; and; (3) at the end, a “station-keeping” `ski` task to make the UUV maintain a fixed position for rendezvous. The survey and station-keeping tasks are IMC plans (discussed in Section III) to be executed by vehicles. Note that the survey task are merely identified by name, thus they are assumed to be pre-programmed for the vehicles, whilst the station-keeping task is programmed inline in the code through the use of an DSL for IMC plans integrated into the Dolphin engine (see Section III-D).

Actual execution of tasks proceeds using an `execute` block. In the code, we see that the `UUVs` are tasked with a composition of three tasks to execute concurrently, as specified by the use of the `|` Dolphin operator: `UUVTask(1) | UUVTask(2) | UUVTask(3)`. Also concurrently, `UUV` is tasked with an `allof` event-based task that works as follows: as each UUV task posts a `ready:i` notification, the `allof` task may consume it in line with guard conditions `when { consume ready: i }` and executing a corresponding `then` block, issuing an IMC plan `rvi` for rendezvous. An `allof` task terminates only when each of the components `when-then` blocks have been completed. In alternative to `allof`, a `oneOf` task, specified with a similar structure, would require only one of the `when-then` blocks to fire, i.e., only one rendezvous would execute instead.

The entire `execute` block terminates when all component tasks terminate. This implies that some vehicles may remain idle (executing some vehicle-dependent fallback behavior such as loitering) while waiting for the completion of ongoing tasks. In sequence, at the end of the program, the `release UUVs + UAV` instruction decouples the vehicles from the program, i.e., `release` is the inverse operation of `pick`. The instruction is redundant in this case, as an implicit release instruction is issued for all bound vehicles at the end of a program, but we show it in the example for the sake of clarity. Note that `pick` and `release` can generally be used at any point in a program to acquire and release vehicles on-the-fly, as illustrated later on (Fig. 5).

C. Task definition

The full definition of Dolphin tasks is summarised by the BNF-style grammar of Fig. 4. In addition to the task operators discussed above, a few others are defined, as follows:

```

Task := PlatformTask // Platform task
| action '{ Code }' // Program-level action
| condition '{ Cond }' // Program-level condition
| Task '>>' Task // Sequential composition
| Task '|' Task // Concurrent composition
| Task '[' VSet ']' // Vehicle set allocation
| allof '{ WhenThen+ }' // All-of block
| oneOf '{ WhenThen+ }' // One-of block (choice)
| waitFor '{ Cond }' // Execution subject
  then Task // to start condition
| until '{ Cond }' // Execution subject
  run Task // to stop condition
| idle Time // Idle task
| during Time // Execution subject
  run Task // to time limit
| watch Task // Error handling
  onError '{ Code }'

WhenThen := when '{ Cond }' then Task
// For IMC-based platform
PlatformTask := imcPlan '{ Id }'
| imcPlan '{ IMC_DSL_Spec }'

```

Fig. 4: Grammar for Dolphin tasks.

— The task allocation operator, `T [V]`, defines the allocation of task `T` to vehicle set `V`, and is already implicit in the example program: `execute v1: T1, v2: T2, ...` (as in the program) is merely syntactic sugar for `execute T1[V1] | T2[V2] | ...`.

— `condition C` prevents progress until condition `C` is satisfied.

— Two operators combine event and control flow, waiting for an condition `C` before starting or stopping a task `T`, i.e., resp. `waitFor C then T` and `until C run T`.

— Two other operators relate to behavior based on a duration of time `t`: idle tasks, `idle t`, or task execution subject to a duration of `t`, `during t run T`.

— Finally, `watch T onError { C }` watches for errors during the execution of `T` (e.g., connection timeouts, internal errors), and executes `C` if one is detected. The code in `C` may include three special actions: `ignore()`, `propagate()`, and `halt()` that respectively ignore the error, propagate it allowing the possibility of being handled by a higher-level `onError` block, or halt the program immediately. A `onError { propagate() }` behavior is the default for each task. When no `onError` block is set to handle an error, the engine halts the program.

```

...
supplier = getSupplier();
while (supplier.active()) {
  // Await for task
  T = supplier.queue.await()
  // Pick UAV
  UAV = pick { type: 'UAV' }
  // Execute
  execute UAV:
    watch
      ( during 10.minutes run T )
    onError {
      err -> {
        message 'Error: ' + err
        ignore()
      }
    }
  // Release UAV
  release UAV
}

```

Fig. 5: Illustration of additional task composition operators.

Some of these additional operators are exemplified in the fragment of Fig. 5. that also provides an illustration of the possible tight integration of Dolphin with Groovy/Java. In the

code, a task supplier is assumed to be implemented externally, possibly mediating interaction with a human user. The supplier yields tasks to be executed by an UAV iteratively, that the Groovy `while` loop executes while the supplier is active. In each iteration a UAV is picked and executes a supplied task for no more than 10 minutes. The constraint is conceivable, e.g., due to battery restrictions. Note that we use one UAV at a time, but different ones may be used in distinct iterations, allowing for vehicle churn in the network. An `onError` block is set notifying the user of any errors with a message, but ignoring it otherwise, thus letting the program resume execution.

D. The Dolphin engine

The Dolphin engine is responsible for executing a program. We now provide some detail of how it works in abstract terms. The main state of a program may be modelled as a partial map representing platform task allocations $\mathcal{A} : \mathcal{V} \mapsto \mathcal{T} \cup \{\perp\}$, where \mathcal{V} and \mathcal{T} are resp. the vehicle and platform task domains and \perp stands for no task allocation, plus a set $\mathcal{B} \subseteq \mathcal{V}$ of vehicles that are bound (associated) to the program. This state varies dynamically according to the use of `pick`, `release`, and `execute`:

- A `pick` block first queries the platform for all connected vehicles \mathcal{V}_c , filters out those in \mathcal{B} (already bound) and those that do not match the selection filters (e.g., type, location, or payload as in the example program of Fig. 3), obtaining a set \mathcal{V}_f . The result \mathcal{V}_p of selected vehicles is a subset of \mathcal{V}_f with a fixed cardinality specified by the `count` parameter (in Fig. 3: 3 in line 6, and 1 by default in line 11) or the full \mathcal{V}_f itself when the `count` parameter is specified as a wildcard value (denoted `_` in the language). The program state is then updated as follows: $\mathcal{B} := \mathcal{B} \cup \mathcal{V}_p$, and $\mathcal{A}[v] := \perp, \forall v \in \mathcal{V}_p$.

- `release` \mathcal{V}_r (as in Fig. 3, l. 40) corresponds to updating the state as $\mathcal{B} := \mathcal{B} - \mathcal{V}_r$, and removing any mappings for \mathcal{V}_r in \mathcal{A} (which should in any case equal \perp at this stage, i.e., no tasks will be executing for \mathcal{V}_r).

- `execute` \mathbf{T} allocates vehicles on-the-fly to platform tasks encoded in \mathbf{T} , i.e., it dynamically changes \mathcal{A} , guided by explicit vehicle allocation specifications through the `[]` operator. If no vehicle allocation is specified, an allowed alternative, the engine allocates all bound vehicles (\mathcal{B}). Given that `execute` only terminates only when the overall \mathbf{T} has fully completed, $\mathcal{A}[v]$ for a vehicle v may alternate between a platform task t , while t executes, and \perp , while v is waiting for task allocation according to global flow of \mathbf{T} . We stick to the informal description of previous sections for the overall flow of composed tasks in terms of sequence, concurrency, or event-flow, operating in the spirit of process calculi [10]. The engine basically executes composed tasks by interpreting the abstract syntax tree of \mathbf{T} and the nature of each operator used in it. A full description is outside the scope of this paper for reasons of space.

In addition to the above characterisation of semantics, event-based program flow through notifications may be accomplished through the `post` and `consume` operations (illustrated in Fig. 3), plus an additional `poll` operation. These

```
public interface Platform ... {
    // Query nodes and tasks
    NodeSet getConnectedNodes();
    PlatformTask getPlatformTask(String id);
    // I/O
    String askForInput(String prompt);
    void displayMessage(String format, Object... args);
    // Extensibility
    void
        customizeGroovyCompilation(CompilerConfiguration cc);
    List<File> getExtensionFiles();
}

public interface Node {
    // Attributes
    String getId();
    String getType();
    Position getPosition();
    Payload getPayload();
    // Task binding & vehicle release
    Task getRunningTask();
    void setRunningTask(Task task);
    void release();
    // Connection handling.
    double getConnectionTimeout();
    void setConnectionTimeout(double timeout); ...
}

public interface Task ... {
    String getId();
    TaskExecutor getExecutor();
    boolean allocate(NodeSet available,
        Map<Task, List<Node>> allocation);
}

public abstract class TaskExecutor ... {
    // Lifecycle methods
    protected abstract void
        onInitialize(Map<Task, List<Node>> allocation);
    protected abstract void onStart();
    protected abstract CompletionState onStep();
    protected abstract void onCompletion(); ...
}

public abstract class PlatformTask implements Task {
    public abstract List<NodeFilter> getRequirements();
    public Optional<Position> getReferencePosition() ...
}
}
```

Fig. 6: Java types for Dolphin platform bindings.

use a simple tuple-space abstraction that can be modelled as multi-set \mathcal{N} of key-value pairs. Each operation works as follows:

- `post` $\mathbf{k}:\mathbf{v}$ adds (\mathbf{k}, \mathbf{v}) to \mathcal{N} .
- `consume` $\mathbf{k}:\mathbf{v}$ removes a (\mathbf{k}, \mathbf{v}) pair from \mathcal{N} , if one exists, and returns true in that case (note that the operation is normally used as a guard condition), otherwise it returns false. If \mathbf{v} equals the wildcard value `_`, the operation tries to remove the oldest tuple in insertion order with key \mathbf{k} .
- `poll` $\mathbf{k}:\mathbf{v}$ merely inspects for the existence of a (\mathbf{k}, \mathbf{v}) pair in \mathcal{N} , also possibly using the `_` wildcard for \mathbf{v} .

E. Platform bindings

For using Dolphin with a concrete platform, a set of Java interfaces and abstract classes are defined for implementation at the platform level. An overview of the main ones are listed in Fig. 6. `Platform` defines the abstract platform operations: querying connected vehicles and existing platform tasks, providing simple user I/O, and extensibility features that include the customisation of Groovy compilation (e.g., automatic imports of certain APIs) and inclusion of DSL extension files to load on startup. Vehicles are instances of `Node`, with associated operations for querying basic attributes (id, type, payload, position), tasking them or releasing them from the program, and connectivity parametrisation. Tasks are instances of `Task`, with an associated id, task allocation

procedure, and task executor instance. A **TaskExecutor** represents a task during execution, defining abstract methods for its lifecycle. Platform tasks are extensions of tasks that must implement **PlatformTask**, providing information regarding vehicle requirements (e.g., vehicle type and payload), and an optional reference position to aid task allocation by the Dolphin engine.

III. INTEGRATION WITH THE LSTS TOOLCHAIN

A. Overview

Over the years, LSTS developed several autonomous vehicles, with an associated open-source toolchain [7]. This toolchain comprises three main components: (1) Neptus, a Java-based command-and-control tool for human operators to configure, plan, and monitor autonomous vehicles using a GUI; (2) DUNE, a C++ on-board software platform for autonomous vehicles, including a simulation mode, and; (3) IMC, an extensible message-based protocol for networked interoperability between all LSTS systems, with bindings in several languages such as C++ and Java.

A subset of IMC is dedicated to the specification, execution, and monitoring of tasks called IMC plans, that we considered as the basic unit of computation for IMC-based Dolphin platforms. An IMC plan is a sequence of maneuvers for a single vehicle, comprising simple maneuvers such as waypoint tracking but also more complex ones such as area surveys. Typically, IMC plans are programmed in Neptus and executed within vehicles by DUNE.

We developed two IMC-based Dolphin platforms taking form as (1) a simple command-line based tool, and (2) integrated in the Neptus tool as a plugin. For both, we had to implement the Dolphin abstract bindings, in particular vehicles and IMC plans respectively as instances of **PlatformTask** and **Node**, presented earlier in Section II-E. In complement, we developed a Groovy DSL specifically devoted to the specification of IMC plans, as we felt the need to define IMC plans directly within Dolphin programs, rather than just relying on the Neptus tool for that purpose.

B. IMC standalone platform

The IMC standalone platform is a command-line tool for executing Dolphin programs. To support IMC interaction, we made use of the IMC Java bindings, and networking code for discovering vehicles in the network using multicast UDP, and then exchanging messages with them over standard UDP. For IMC plan interaction, we made use of **PlanSpecification** (for plan definition), **PlanControl** (control), and **PlanControlState** (monitoring) messages from the IMC specification². The platform extends the Dolphin DSL with **imcPlan** tasks, identified by id or defined inline using the IMC DSL (discussed below).

C. Dolphin plugin for Neptus

The Dolphin plugin for Neptus allows users to edit and run Dolphin programs through a custom window, embedded in the overall GUI environment for editing and monitoring




Fig. 7: Dolphin plugin running in Neptus.

```
t = imcPlan {
  // Id
  planName 'waterSurvey'
  // Set reference speed, depth, and location
  speed 1.5, Speed.Units.METERS_PS
  z      0.0, Z.Units.DEPTH
  locate Location.APDL
  // Goto maneuver, activating the camera payload
  move 30, -125
  goTo payload: [[name: 'Camera']]
  // Loiter maneuver
  move (-30, -50)
  loiter radius:100
}
```

Fig. 8: Task specified using the IMC DSL.

of IMC plans, as illustrated in Fig 7. During execution, the behavior of a program can be simultaneously monitored using the Dolphin console and the standard Neptus GUI. The implementation traits are similar to the stand-alone IMC platform, apart from a delegation of networking functions to the pre-existing Neptus infrastructure, and an integration with the database of IMC plans associated to a Neptus console.

D. The IMC DSL

The IMC DSL may be used to automate IMC plan generation with relatively succinct textual descriptions. The listing of Fig. 8 illustrates a richer example than the one given in the sample program of Fig. 3, defining an IMC plan with two maneuvers (**Goto** and a **Loiter**), and associated parameterisation.

IV. FIELD TESTS

Dolphin has been evaluated in field tests that took place at the Leixões harbour, and in open sea at Tróia during the 2017 Rapid Environmental Picture (REP'17) exercise³ in collaboration with the Portuguese Navy [16, 17]. Here we only present results for the example scenario/program of Section II-B at Leixões.

A. Setup

In Fig. 9 we present an overview of the scenario (9a) and a photo of the vehicles we used (9b). The scenario overview is a Neptus screenshot, with the scenario already being executed, depicting three IMC plans for the water surveys to be executed by UUVs, and one of the rendezvous plans executed by a simulated UAV (the circular loiter). Due to operational restrictions at Leixões, we could only deploy

²<http://github.com/LSTS/imc>

³<http://rep17.lsts.pt>



(a) IMC plans in Neptus.




(b) Noptilus UUVs (orange-colored vehicles).

Fig. 9: Test setup.

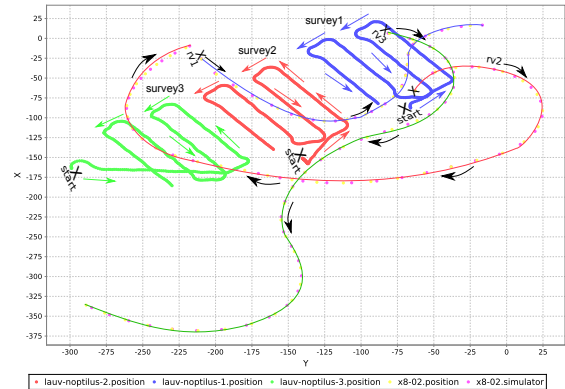
a simulated UAV as part of the control loop. The harbour location is shown bottom right in the same image⁴. For the tests, we used three LAUV-class vehicles [18] shown in the photo of Fig. 9b (an additional fourth vehicle shown was used for unrelated operations), named Noptilus-1, Noptilus-2, and Noptilus-3. The same photo shows a Manta communications gateway [7], used for WiFi and underwater communications between vehicles and Neptus consoles.

Each of the UUVs was equipped with a DVL, allowing us to fulfil the purpose of gathering bathymetry data (distance to the seafloor) for the area of operation, thus the vehicles were programmed to operate at the surface. A variant of the program in Fig. 3 was edited in Neptus using the Dolphin plugin, and the associated IMC plans were edited using a Neptus plan edition console. The main differences of the program variant were that we used named vehicle selection (through an `id` attribute in `pick` blocks) to avoid unexpected movements in the relatively short area we had for operation, and that the IMC plans for UUVs did not include a station-keeping maneuver after the survey. Moreover, we left the connection timeout for UUVs (through `setConnectionTimeout`, discussed in Section II) to the default minimum of 1 minute, since they operated at the surface and always in WiFi reach. During tests in Tróia [16, 17], the same overall scenario was executed, but a timeout of several minutes was set instead to cope with connectivity issues raised by the use of underwater acoustic communications. The difference is that, in Tróia, the UUVs operated underwater for long periods with a bottom-tracking approach (i.e., maintaining a constant distance to the seafloor) for the

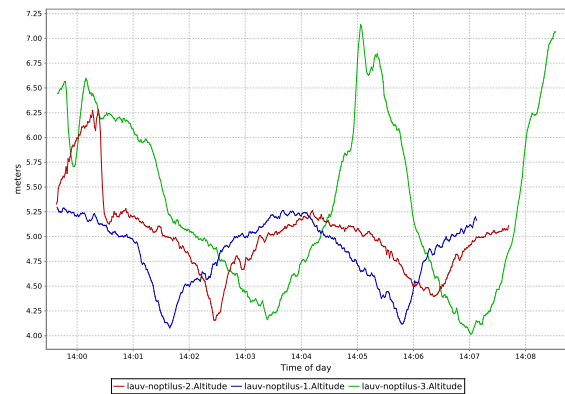
⁴<https://goo.gl/maps/hHvdiTt3xAH2>




(a) Execution timeline.



(b) Vehicle positions (XY).



(c) Bathymetry values measured using DVL.



(d) Bathymetry map.

Fig. 10: Results for the field test scenario.

purpose of gathering side-scan sonar data.

B. Results

The results of one of the executions of the target scenario are shown in Fig. 10a comprising: a timeline of executed IMC plans (10a), vehicle positions (10b), DVL measurements over time per vehicle (10c), and a global bathymetry plot derived from the DVL measurements (10d).

As shown in the timeline of Fig. 10a, the whole execution took roughly 9 minutes. The UUVs started their surveys

simultaneously, but each of the three surveys terminated at different times. Even if the surveys had similar path lengths, non-linear factors such as initial position, sea currents, vehicle calibration obviously impact on execution times. As each survey terminated, the UUV initiated a corresponding rendezvous. Purely by chance, the order of completion of the survey and rendezvous maneuvers was in line with vehicle numbering (1, 2, 3). Other executions of the scenario yielded a different order of completion.

The XY plot in Fig. 10b shows that the UUV paths are in line with the programmed survey plans. In the same figure, for clarity, the UUV path is annotated in terms of rendezvous plan trajectories (**rv1**, **rv2**, and **rv3**), where we can notice an approach of the UAV to each survey area. For a non-cluttered plot, we omit the UAV paths in between rendezvous plans, during which the vehicle loitered in the air. The DVL values of Fig. 10c indicate different depths and variations according to each vehicle/survey area, with values ranging approximately from 4.0 to 7.5 meters. The vehicle logs were collected and post-processed to obtain the bathymetry plot of the overall operation area in Fig. 10d.

V. RELATED WORK

Dolphin follows on from NVL [15], also a task orchestration language. Similarly to Dolphin, NVL defines primitives for selecting vehicles over a network, but only a single “step” primitive for firing tasks concurrently for multiple vehicles in contrast to compositionally-defined tasks in Dolphin. NVL requires explicit task-vehicle allocation with the granularity of single vehicles rather than vehicle sets, and is not extensible out-of-the-box for new constructs or robotic platforms, as these require direct changes on the base code of NVL. Integration with the LSTS toolchain is also very limited: apart from the use of IMC, a command line executor and a language-specific editor had to be used with no interface to Neptus. It also relied on previously programmed IMC plans, in contrast to the Dolphin integration where we have the option of using the IMC DSL.

We now survey other DSLs for coordinated task execution of networked robotic systems, then make a final discussion contextualising Dolphin in the overall research landscape.

Karma [19] implements an orchestration architecture for programming micro-aerial vehicle (MAVs) networks, called the hive-drone model. Orchestrated behavior is conducted by a centralised coordinator, called the hive. Tasks, called drone behaviours, can be allocated to multiple vehicles and interact through a centralised datastore running at the hive. Each behavior is specified independently, with an associated activation predicate and a progress function, both of which feed on information the hive datastore. The progress function governs the on-the-fly allocation of more or less drones by the hive to a single behavior, according to task completion results reported by drones. Though there are no composition constructs, composed behaviours may be defined implicitly by datastore value dependencies. The hive-drone model of Karma is also used by the Simbeotic simulator for MAV swarms [20]. Other languages used with aerial vehicles

follow the spirit of Karma’s centralised architecture, using different kinds of abstractions, e.g., TeCoLa [21] is a Python DSL built around the notions of vehicle teams and services accessible via remote procedure calls, and CSL [22] uses reconfigurable Petri nets for task orchestration.

Proto [23] is a functional programming language for homogeneous robots. The conceptual approach is based on computational fields, whereby a collection of devices approximates a continuous field in space/time. A Proto program specifies choreographed swarm-like behavior through the composition of operators for restricting execution in space and time, feedback-loops that define state and execution flow, and neighbourhood-based computation. Programs are compiled to abstract bytecode that is deployed using a viral propagation mechanism over the network, and then executed by each robot in distributed manner. Bytecode execution uses a stack-based virtual machine for programs that may run on very lightweight microprocessor chips with only a few KB of RAM. Protelis [24] is a more recent language based on Proto, embedded in/interoperable with Java.

Meld [25] is a logic programming language that, like Proto, also realizes a top-down synthesis approach. The state at each robot defined by a set of logical facts that evolve according to rules producing new facts. Rules take into account local robot state, but also a special constructs that access the state of neighbouring robots. Like Proto, the Meld compiler derives local robot programs that run on very lightweight embedded platforms.

Buzz [26] is a DSL for programming heterogeneous robot swarms. Programs directly run locally on each robot, instead of being derived from top-level specification as in Proto or Meld, but have a notion of belonging to a specific swarm with the intent of emergent collective behavior. Multiple swarms of heterogeneous robots may be formed, making use of primitives for swarm formation, neighbourhood-based queries and broadcast operations, and additional communication through a distributed tuple space. Buzz is compiled onto abstract bytecode that is executed by a lightweight virtual machine, and can be integrated with C/C++ code. In particular, ROSBuzz [27] integrates Buzz in ROS. In the line of Buzz, Swarmorph-script [28] is a rule-based language for self-assembling robot swarms for morphogenesis.

Voltron [29] is a language for mobile sensing using autonomous vehicles. The network of available vehicles is tasked as a whole, regardless of how many vehicles are available and without having to associate tasks to vehicles. Tasks are defined by actions to be executed at a set of locations, using a key-value store for coordinated behavior, and can be started/stopped and engage less/more vehicles on-the-fly, in line with an active sensing strategy that accounts for the evolution of accomplished goals. Voltron is implemented through source-to-source translators to C++ and Java, and supports centralised and distributed execution modes. In the distributed execution mode, every vehicle bound to the same task executes the same program (as in Proto), and virtual synchrony mechanisms are employed for the consistency of the shared key-value store.

Summarising the above discussion, we observe three main types of approach for globally tasking autonomous vehicle networks: (1) orchestration, where a single program/coordinator tasks nodes on-the-fly without need for explicit coupling between nodes and optional node-to-node communication (Dolphin, NVL, Karma, TeCoLa, CSL); (2) choreography, where nodes are programmed through a global specification (Proto, Protelis, Meld) leading to synthesised programs that organise as swarms, and; (3) distributed programs without an explicit global specification but emergent swarm behavior (Buzz, Swarmorph-script). Voltron allows both orchestration and choreography, given the choice between centralised and distributed implementations.

Dolphin shares common features to the discussed languages, such as: the ability of tasking vehicle teams in most of them; a design for extensibility and integration with other languages as in the case of TeCoLa, Buzz, Protelis, or Voltron, and; an explicitly compositional definition of tasks as in Proto and Meld. In regard to the later aspect, a distinguishing feature of Dolphin lies in the use of a process-calculi approach for task definition. On the other hand, Dolphin lacks relevant features, some of which discussed as future work in the next section: tasks that are associated dynamically to multiple vehicles in line with a notion of progress as in Karma and Voltron, or the possibility of neighbour-based and team-level primitives for cooperative behavior as in Meld, Proto, Buzz, and Voltron.

VI. CONCLUSION

We presented Dolphin, a programming language for task orchestration in autonomous vehicle networks, its integration with the LSTS toolchain for autonomous vehicles, and the use of the language in a field tests involving multiple vehicles. A Dolphin program is a global specifications of multi-vehicle tasks that are defined compositionally. The language is also naturally extensible by virtue of its definition as a Groovy DSL and of the abstract platform bindings that made the LSTS toolchain integration possible.

As future work, we are interested in extending Dolphin in a number of ways such as: the representation of human operators or sensors as nodes, in addition to vehicles; vehicle interaction constructs in support of cooperative tasks, in complement to the centralised tuple-space scheme we now employ; tasks that aggregate vehicle sets with varying cardinality, for example according to progress measured as set of accomplished goals or in reaction to vehicle faults during execution, and; more expressive operators for a space-time characterisation of task flow, for now only implicit in vehicle selection criteria or the nature of IMC-based plans we used in the LSTS platform. Additional platform bindings for popular toolkits such as ROS [5] would also be interesting, beyond the current support for the LSTS toolchain and the work in progress regarding MAVLink [6, 9].

REFERENCES

[1] M. Dunbabin and L. Marques, "Robots for environmental monitoring: Significant advancements and applications," *IEEE Robotics Automation Magazine*, vol. 19, no. 1, pp. 24–39, 2012.

[2] R. B. Wynn, *et al.*, "Autonomous Underwater Vehicles (AUVs): Their past, present and future contributions to the advancement of marine geoscience," *Marine Geology*, vol. 352, pp. 451 – 468, 2014.

[3] C. Petrioli, *et al.*, "The SUNSET framework for simulation, emulation and at-sea testing of underwater wireless sensor networks," *Ad Hoc Networks and Physical Communication*, pp. 224 – 238, 2015.

[4] I. R. Castro and M. Parashar, "Architecting the cyberinfrastructure for National Science Foundation Ocean Observatories Initiative (OOI)," in *Instrumentation Viewpoint*, no. 19. SARTI, 2016, pp. 99–101.

[5] M. Quigley, *et al.*, "ROS: an open-source robot operating system," in *Proc. ICRA'09*, vol. 3, no. 3.2. IEEE, 2009, p. 5.

[6] "MAVLink: Micro air vehicle communication protocol," <https://mavlink.io/en/>.

[7] J. Pinto, *et al.*, "The LSTS toolchain for networked vehicle systems," in *Proc. MTS/IEEE OCEANS'13*. IEEE, 2013, pp. 1–9.

[8] M. Gombolay, *et al.*, "Computational design of mixed-initiative human-robot teaming that considers human factors: situational awareness, workload, and workflow preferences," *International Journal of Robotics Research*, vol. 36, no. 5–7, pp. 597–617, 2017.

[9] "Dolphin website," <https://DolphinDSL.github.io>.

[10] R. Milner, "A calculus of communicating systems," *LNCS*, vol. 92, 1980.

[11] F. Dearle, *Groovy for Domain-Specific Languages*. Packt Publishing, 2015.

[12] L. L. Sousa, *et al.*, "Integrated Monitoring of Mola Mola Behaviour in Space and Time," *PLOS one*, vol. 11, no. 8, p. e0160404, 2016.

[13] J. Pinto, *et al.*, "Chasing fish: Tracking and control in a autonomous multi-vehicle real-world experiment," in *Proc. Oceans'13*. MT-S/IEEE, 2013.

[14] M. Faria, *et al.*, "Coordinating UAVs and AUVs for oceanographic field experiments: Challenges and lessons learned," in *Proc. ICRA'14*. IEEE, 2014, pp. 6606–6611.

[15] E. R. B. Marques, *et al.*, "NVL: a coordination language for unmanned vehicle networks," in *Proc. SAC'15*. ACM, 2015, pp. 331–334.

[16] K. Lima, *et al.*, "Programming networked vehicle systems using Dolphin – field tests at REP'17," in *Proc. IEEE/MTS Oceans'18*. IEEE, 2018.

[17] K. Lima, "Dolphin: A Domain-Specific Language for Autonomous Vehicle Networks," Master's thesis, MIERSI/FUCUP, 2017.

[18] L. Madureira, *et al.*, "The Light Autonomous Underwater Vehicle: Evolutions and Networking," in *Proc. MTS/IEEE OCEANS'13*. IEEE, 2013, pp. 1–6.

[19] K. Dantu, *et al.*, "Programming micro-aerial vehicle swarms with Karma," in *Proc. SenSys '11*. ACM, 2011, pp. 121–134.

[20] B. Kate, *et al.*, "Simbeotic: a simulator and testbed for micro-aerial vehicle swarm experiments," in *Proc. IPSN'12*. ACM, 2012, pp. 49–60.

[21] M. Koutsoubelias and S. Lalis, "TeCoLa: A programming framework for dynamic and heterogeneous robotic teams," in *Proc. MOBIQUITOUS'16*. ACM, 2016, pp. 115–124.

[22] E. Pereira, *et al.*, "The C3UV testbed for collaborative control and information acquisition using UAVs," in *Proc. ACC'13*. IEEE, 2013, pp. 1466–1471.

[23] J. Bachrach, *et al.*, "Composable continuous-space programs for robotic swarms," *Neural Computing and Applications*, vol. 19, no. 6, pp. 825–847, 2010.

[24] D. Pianini, *et al.*, "Protelis: Practical aggregate programming," in *Proc. SAC'15*. ACM, 2015, pp. 1846–1853.

[25] M. P. Ashley-Rollman, *et al.*, "A language for large ensembles of independently executing nodes," in *Proc. ICLP'09*. Springer, 2009, pp. 265–280.

[26] C. Pinciroli and G. Beltrame, "Buzz: An extensible programming language for heterogeneous swarm robotics," in *Proc. IROS'16*. IEEE, 2016, pp. 3794–3800.

[27] D. St-Onge, *et al.*, "ROS and Buzz: consensus-based behaviors for heterogeneous teams," *arXiv preprint arXiv:1710.08843*, 2017.

[28] R. O'Grady, *et al.*, "Swarmorph: Morphogenesis with self-assembling robots," in *Morphogenetic Engineering*. Springer, 2012, pp. 27–60.

[29] L. Mottola, *et al.*, "Team-level programming of drone sensor networks," in *Proc. SenSys '14*. ACM, 2014, pp. 177–190.

Acknowledgements This work was partially funded by European Commission, ERDF (P2020), under the projects Endurance (NORTE-01-0247-FEDER- 017804) and SMILES/TEC4GROWTH (NORTE-01-0145-FEDER-000020).