

SwarmScript 3.0

Algorithmic Specification

Sebastian von Mammen, Scott Steil, and Christian Jacob

University of Calgary, Canada

1 SwarmScript—A High-level Description

SwarmScript provides a powerful programming approach based on situation-action pairs which allow the modeller to express agent behaviours and which equally well translate into computational processes. It merges a standardized representation for modelling agent behaviours with an algorithmic interpretation for simulating the resulting interactions.

In SwarmScript, situation-action pairs are built bottom-up from selection and action operators. The basic idea is that selection operators pass situational information to actions which change the simulation state. The information flow is determined through connections that are hooked up to plugs exposed by the operators. Information about the simulation state flows through chains of selection operators that feed into action operators.

On the one hand, chains of selection operators serve as sources of information for an action, e.g. they trigger whether or not an action should be executed and provide it with various parameters. These selection operators are classified as *source* selections. On the other hand, chains of selection operators direct an action toward specific targets. These selection operators are classified as *target* selections. Given that sources and targets are clearly separated, sources encompass all the information that motivates and informs an action, whereas targets clearly denote those programmatic objects that are changed based on the execution of an action. This distinction highlights the relationship between cause and effect, the two fundamental, interwoven constituents of an interaction.

The next section will introduce the data model of SwarmScript v3.0, whereas Sections 3 to 5 present pseudo code for creating, modifying and running SwarmScript simulations. Section 3 introduces some technical methods for setting, updating, and querying the connectivity of SwarmScript operators. Section 4 sheds light on routines necessary for creating and editing SwarmScript behaviours, including the creation of high-level operators. Finally, the pseudo code in Section 5 shows how SwarmScript behaviours can be simulated.

Some methods and functions that occur in the presented pseudo code are typically provided by high-level programming libraries and either access the properties of abstract data structures (getters/setters and type checking), create or delete instances (constructors receive their properties as parameters in the order shown in Figure 1), or they are generic operators that work on collections. In particular, the following list-related routines are assumed to be commonly understood and available without further explanations:

- *list.append(element)*, *list.append(list)*,
- *list.remove(element)*,
- *list.hasNext()* and *list.next()*,
- *list.hasElement(element)*, *list.hasElements(list)*,
- *list.first()*, *list.last()*, *list.get(index)*,
- *list.size()*,
- *list.copy()*, and
- *list.sort()*.

Regarding the usage of basic control flow routines, *for* and *while* loops are used whenever the iteration order is important. Otherwise, *for each* statements are deployed. *For all* is only used once, in a case in which the most efficient way to retrieve the operands depends on the implementation framework.

2 Language Elements

In this section, the role of the various elements of the SwarmScript language is presented in detail and an according data model is introduced. Figure 1 shows a schematic diagram of the involved object classes and their dependencies in UML/Martin notation (Unified Modelling Language). Different Operator subclasses, i.e. Selections, Actions, and BehaviourFrames, have Plugs that are interconnected by means of Connection instances.

Figure 2 depicts a schematic visualization of a potential agent behaviour expressed by intertwined SwarmScript language elements. Connections are represented as thin black lines, small circles denote plugs, the wrapping boxes are operators, small numbers in the upper right corner of an operator indicate its hierarchical level, the *ptrigger* plug of a behaviour frame is depicted as an on/off switch in the upper left corner of a behaviour frame. A hierarchical operator can be expanded to reveal its underlying operators. In such an open state, the higher-level operator is reduced to a framing rectangle enclosing its underlying operators with a button in the upper right corner. This button indicates that it can be wrapped, reduced to its compact representation at any point in time.

Actions and behaviour frames are placed on the grey strip, whereas source selections are arranged on the left-hand side and target selections on the right-hand side of the strip. Information flows through chains of selection operators from their output plugs to the connected input plugs.

Actions and behaviour frames only have input plugs. They receive information from source selections from the left and from target selections from the right. The spatial separation of selections is in accordance with their semantics is also reflected in the arrangement of input plugs of actions and behaviour frames—source input comes in from the left, target input from the right.

The following paragraphs detail the concepts underlying the individual language elements and how they work together.

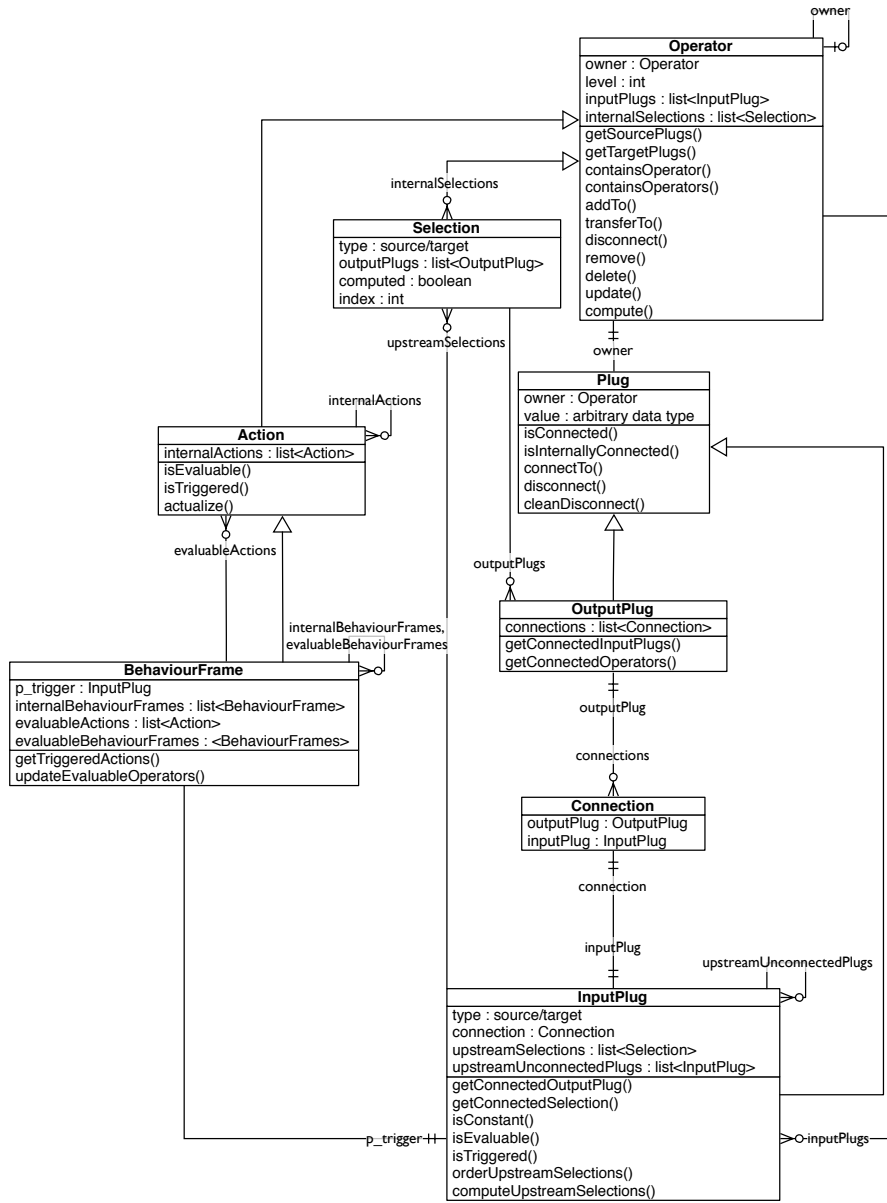


Fig. 1. The classes of the SwarmScript language elements in UML/Martin notation.

2.1 Connection

Connection instances link two operators by referencing one **input** and one **out-put plug**, respectively.




Fig. 2. A schematic visualization of SwarmScript language elements as they might be intertwined to express agent behaviour. The red arrows and labels denote the translation of the visuals into SwarmScript elements.

2.2 Plug

A Plug instance belongs to an operator, its **owner**. Additionally, it has a value assigned—by default it is *nil*.

2.3 OutputPlug

Output plugs can share any number of **connections** to input plugs.

2.4 InputPlug

Based on the intended use of the information received from the selection operators, input plugs are considered either of **type** source or target. An input plug may have one **connection** to an output plug. The sequence of selections that feeds into an input plug is stored in **upstreamSelections**, unconnected input plugs that occur in this sequence are stored in **upstreamUnconnectedPlugs**.

2.5 Operator

Selection, Action, and BehaviourFrame are directly or indirectly derived from the Operator class. An operator may be owned by another operator, its **owner**. The level of recursion of an operator is reflected by its **level** property. It is determined by the maximal recursion level of its underlying operators plus one. Any operator has a (possibly empty) list of input plugs (**inputPlugs**) which provide information to the operator, and a list of Selection instances which provide or refine the incoming information. An operator can be built hierarchically and host multiple underlying selection operators. As there might be dependencies regarding the execution order, these selections are stored in the ordered list of **internalSelections**. Unconnected input and output plugs of the internal operators become the input and output plugs of the higher level operator.

2.6 Selection

Selection operators provide situational information that informs or triggers, i.e. instigates, an agent's interactions. Selection operators can be joined into chains by connecting their output plugs to other operators' input plugs. Although the propagation of information across chains of selection operators can result in complex computations, selections do not change the simulation. Selections only pass information to actions which in turn change the simulation state based on the provided information.

A selection operator stores values in its output plugs based on information received from its input. It can, for example, filter a list of agents by their name (one input and one output plug) or merge two lists of agents and pass it on (two input plugs, one output plug). Selection operators without inputs can output information that is independent of preceding selection operators. Such terminal selection operators could, for instance, provide a constant number, the outcome of a random experiment, or forward information retrieved from external computational engines, e.g. objects that have collided with a given agent determined by a physics engine.

If a selection operator fails to compute, for instance because it is lacking important information, it **might** store a *nil* value in its output plug(s). An action operator that receives one or more *nil* values, will not be executed.

There are two **types** of selections. Source selections provide information to trigger and/or configure an action, whereas target selections identify those programmatic objects that will be changed when the action is executed. This distinction is mainly semantic at the model level but might be reflected at the user-interface level as shown in Figure 2.

Individual selection operators might feed their information into multiple action operators—either because the actions target the same programmatic objects or because they retrieve information from the same sources. Instead of evaluating shared selection operators multiple times, they are computed only once per iteration and their **computed** properties are set to true.

A selection's **index** property is utilized when it is sorted to resolve the execution interdependencies among operators.

2.7 Action

An action operator only has input plugs. These input plugs are classified as source and target plugs in accordance with their intended usage. An action retrieves information from its source plugs and applies changes to selections connected to its target plugs.

If one or more input plugs of a connected selection chain are unconnected or set to *nil*, an action is not evaluable, and thus, not considered for execution. All evaluable actions, however, are considered for execution at each simulation step. If their associated selection operators provide valid return values (non-*nil*), the actions get actualized.

An action can be comprised of a sequence of several underlying actions. Naturally, the whole sequence is only executed, if (1) all of the underlying actions are evaluable, and (2) all of the associated selections provide valid return values.

An action or an action sequence can be wrapped into high-level actions together with (parts of) the connected selection chains. Unconnected source or target plugs of the wrapped **internalActions** are adopted by the new operator as such, whereas unconnected input plugs of the associated selections are added to its source or target plugs depending on their types.

2.8 BehaviourFrame

Behaviour frames can wrap sets¹ of **internal** selections, actions, and **behaviour frames**. Similar to actions, behaviour frames only expose input plugs of types source and target. However, different from actions, the exposed plugs of a behaviour frame do not need to be connected in order to consider the encapsulated actions for execution. Any evaluable, encapsulated actions are still considered for execution.

However, the execution of a behaviour frame may be skipped altogether, in case its activation plug **p_{trigger}** is not evaluable². Therefore, as a first optimization step, evaluable actions and lower-level behaviour frames whose *p_{trigger}* plugs are evaluable are stored in a behaviour frame's **evaluableActions** and **evaluableBehaviourFrames** lists. In a second step, only these evaluable behaviour frames and actions will be considered for execution.

2.9 Schematic Instance Representation

The series of Figures 3 to 7 depicts the hierarchical design of a possible SwarmScript behaviour. In Figure 3, a sequence of two action operators is shown at

¹ In numerous cases, *sets* would be mathematically more adequate as containers for SwarmScript elements than *lists*. However, due to their topological relationships, it is important to store most SwarmScript operators in ordered *lists*. Using lists facilitates the maintenance of an expedient, consistent relationship between the syntactical arrangement and the semantical interpretation of their symbolical representation.

² Due to the design of SwarmScript, especially because actions only execute, if their input plugs have or receive valid information, the return value *nil* generally implies no execution.

the top of the grey strip, followed by another single action in the centre, and a behaviour frame at the bottom.

The sequence of two action operators at the top will execute in top-down order, only if (a) both operators are fully connected and (b) all their plugs receive non-*nil* values. Due to their interdependency, sequences of actions can be wrapped into a higher order action operator. As one can see, the given action sequence will not be considered for execution, since its second plug on the left-hand side is not set.

The action operator in the centre of the grey strip, it seems, is fully connected. However, upon closer inspection, one will find that the left-most selection operator, which this action operator also depends on, is not fully connected. Therefore, this action operator, too, will not be considered for execution in the given configuration. The same is true for behaviour frame operator at the bottom of the grey strip³.

In Figure 4, the level-2 selection operator to the left is opened up for inspection and modification. It shows that its output is generated by passing information through four internal selection operators, one of which is hierarchically designed as well (level-1). The second action of the action sequence at the top of the grey strip is opened up in Figure 5, revealing that it consists of another sequence of two action operators. In Figure 6, the behaviour frame at the bottom of the grey strip is opened. It consists of two behaviour frames itself and four associated selection operators. Diving into the bottom-most behaviour frame one more time in Figure 7 reveals the combination of one action sequence and one further, independent action, depending on another nine internal selection operators.

3 Basic Routines

This section introduces routines that revolve around the connectivity, i.e. retrieving, setting, and changing the state of connectivity of SwarmScript operators. These rather technical routines are fundamental for creating and editing a SwarmScript behaviour (Section 4) and for processing it as part of a simulation (Section 5).

3.1 Additional Getter Methods

By means of their respective connections, plugs can be queried for their linked plugs or linked operators. In particular, an `OutputPlug` instance can be queried

³ The behaviour frame is visualized with a black bar at its centre, which not only resembles the grey strip of the behaviour frame that hosts the complete behaviour of an agent. But it also indicates that the source information on the left-hand side of the operator might not be related to the changes induced to the targets connected on the right-hand side. In actions, on the other hand, the internal logic between sources and targets is intertwined to such a degree that they can be considered completely dependent. Therefore, actions are not shown with a bar between the two kinds of input plugs.




Fig. 3. A possible SwarmScript behaviour in the works—all actions and the depicted behaviour frame are not fully connected, yet.




Fig. 4. The selection operator at the left centre has been opened up for inspection/modification.




Fig. 5. The second action operator of the action sequence on the grey strip is itself composed of a sequence of two action operators.




Fig. 6. The behaviour frame at the bottom consists of two lower-level behaviour frames and four selection operators.




Fig. 7. Nine more selection operators, an action sequence and another independent action were hidden in the bottom-most behaviour frame.

for its linked input plugs (**getConnectedInputPlugs()**, Alg. 1) or linked operators (**getConnectedOperators()**, Alg. 2). An InputPlug instance can be queried for its associated output plug (**getConnectedOutputPlug()**, Alg. 3) or selection operator (**getConnectedSelection()**, Alg. 4).

Output: *list*(InputPlug)
plugs_{connected} := {};
for *i* := 0, *i* < *self.connections.size()*, *i* ++ **do**
plugs_{connected}.append(self.connections.get(i).inputPlug);
return *plugs_{connected}*;
Algorithm 1: *OutputPlug.getConnectedInputPlugs()*

Output: *list*(Operator)
operators_{connected} := {};
plugs_{connected} := *self.getConnectedInputPlugs()*;
for *i* := 0, *i* < *plugs_{connected}.size()*, *i* ++ **do**
operators_{connected}.append(plugs_{connected}.get(i).owner);
return *operators_{connected}*;
Algorithm 2: *OutputPlug.getConnectedOperators()*

Output: OutputPlug
return *self.connection.outputPlug*;
Algorithm 3: *InputPlug.getConnectedOutputPlug()*

Output: Selection
connectedOutputPlug := *self.getConnectedOutputPlug()*;
if *connectedOutputPlug == nil* **then**
| **return** *nil*
end
else
| **return** *connectedOutputPlug.owner*;
end
Algorithm 4: *InputPlug.getConnectedSelection()*

It is important to know which action plugs and *p_{trigger}* plugs a specific input plug affects based on the given connection topology. Any topological changes

require these plugs to update their dependencies. **getAffectedInputPlugs()** (Alg. 5) finds an according list of affected, non-Selection plugs. The function tests whether the given input plug belongs to a non-Selection operator, otherwise it recursively steps through all the output plugs of its owner and recursively considers their linked input plugs. Figure 8 depicts the algorithmic process effected by the *getAffectedInputPlugs()* routine. Starting with the input plug of the left-most selection operator, it forward traverses all linked selection operators until it connects to some of the action operators located on the grey strip, to the *p_{trigger}* plug of the outer, opened behaviour frame, and to one of its selection’s input plugs, which is exposed as one of the behaviour frame’s exposed input plugs.

The latter case deserves some attention: When building the hierarchical behaviour frame operators, the Selection instance’s input plug became a source plug of the inner behaviour frame first, and of outer behaviour frame afterward. Therefore, it is not considered a plug of a Selection instance any longer. As high-level operators become the owners of their exposed plugs (implemented by the updating routines, Algs. 36 and 38), the internal selection operators of the opened selection operator are not traversed.

```

Output: list(InputPlug)
if class of self.owner is not Selection then return {self};
downstreamAffectedInputPlugs := {};
foreach pout ∈ self.owner.outputPlugs do
    foreach pin ∈ pout.getConnectedInputPlugs() do
        | downstreamAffectedInputPlugs.append(pin.getAffectedInputPlugs());
    end
end
return downstreamAffectedInputPlugs;
Algorithm 5: InputPlug.getAffectedInputPlugs()

```

The Operator base class provides input and output plugs and according, inferable setter and getter methods. The input plugs of action operators and behaviour frames are divided into source and target plugs. The Operator methods **getSourcePlugs()** (Alg. 6) and **getTargetPlugs()** (Alg. 7) provide easy access to these two distinct types of input plugs.

3.2 State Querying Routines

A plug can either be connected or not. The according state querying function **isConnected()** tests whether the list of connections of an output plug is initialized and not empty (Alg. 12), its implementation for an input plug simply tests whether or not its connection reference is initialized at all (Alg. 13).

Except for a “root” behaviour frame which contains all of an agent’s SwarmScript behaviour, any operator is contained inside another one. Some plugs of an

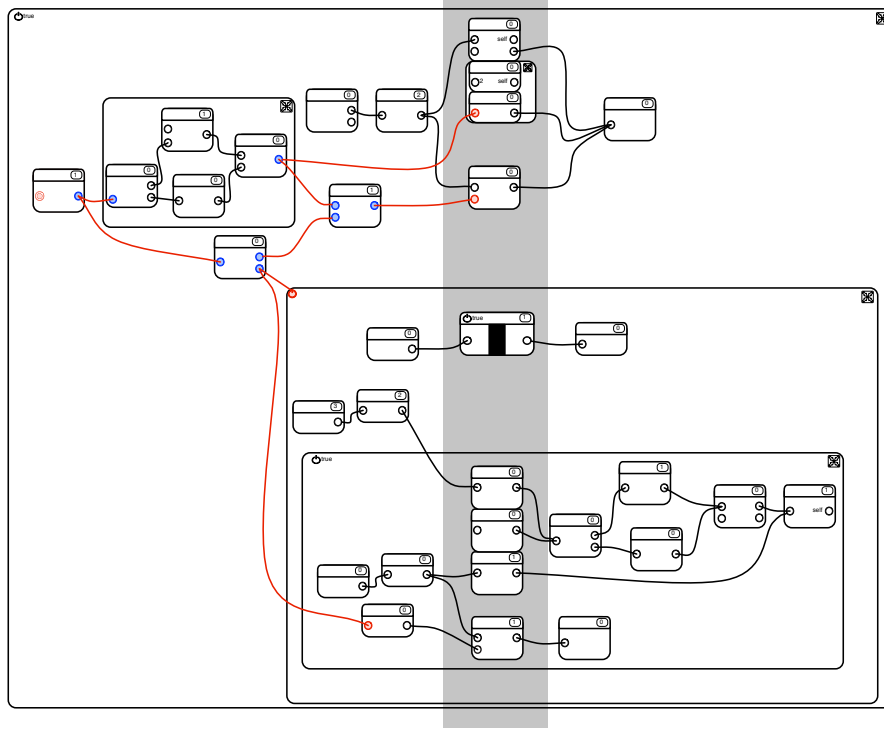


Fig. 8. Starting from the left-most selection operator, all affected non-Selection input plugs are discovered by the function *getAffectedInputPlugs()*. The result are the red coloured plugs of operators on the grey strip as well as the single source and the $p_{trigger}$ plug of the opened, outer behaviour frame.

Output: *list*(*InputPlug*)
 plugs := {};
 for $i := 0, i < self.inputPlugs.size(), i++$ do
 | if $self.inputPlugs.get(i).type == source$ then
 | | plugs.append($self.inputPlugs.get(i)$);
 end
 return plugs;

Algorithm 6: *Operator.getSourcePlugs()*

Output: *list*(*Plug*)
 plugs := {};
 for $i := 0, i < self.inputPlugs.size(), i++$ do
 | if $self.inputPlugs.get(i).type == target$ then
 | | plugs.append($self.inputPlugs.get(i)$);
 end
 return plugs;

Algorithm 7: *Operator.getTargetPlugs()*

Input: ops : $\langle \text{Operator} \rangle$
foreach *Operator* $o \in \text{ops}$ **do**
 | **if** *self.containsOperator*(o) == *false* **then return false;**
end
return true;

Algorithm 8: *Operator.containsOperators()*

Input: o : *Operator*
if *class of o is not Selection* **then return false;**
**return self.internalSelections.hasElement(o);
Algorithm 9: *Selection.containsOperator()***

Input: o : *Operator*
if *class of o is Selection* **then**
 | **return self.internalSelections.hasElement(o);
else if *class of o is Action* **then**
 | **return self.internalActions.hasElement(o);
end
return false;
Algorithm 10: *Action.containsOperator()*****

Input: o : *Operator*
if *class of o is Selection* **then**
 | **return self.internalSelections.hasElement(o);
else if *class of o is Action* **then**
 | **return self.internalActions.hasElement(o);
end
**return self.internalBehaviourFrames.hasElement(o);
Algorithm 11: *BehaviourFrame.containsOperator()*******

Output: boolean
return self.connections != nil AND self.connections.size() > 0;
Algorithm 12: *OutputPlug.isConnected()*

Output: boolean
return self.connection != nil;
Algorithm 13: *InputPlug.isConnected()*

operator might be connected from the outside, i.e. they might be connected to operators that are not internal operators of their operator's owner. These plugs are usually part of the exposed plugs of the higher level operator. However, at the time of the creation of the high-level operator, they might only be identified by checking whether or not the operators of their connected plugs are part of the same subtree in the operator hierarchy or not. The according test first queries whether a plug is connected at all, and second whether its connected operator(s) are all at the same level as itself, or not (**isInternallyConnected()**, Algs. 14 and 15).

Output: boolean
return self.isConnected() AND
self.owner.owner.containsOperators(p.getConnectedOperators());
Algorithm 14: *OutputPlug.isInternallyConnected()*

While the value of an output plug is updated as part of the execution of the internal logic of the respective operator, input plugs pull in information, allowing for its propagation across chained operators. It is important to know whether it is possible for an input plug to receive a valid value or not—an invalid value (*nil*) at a non-selection input plug prevents its operator from execution. We say an input plug is evaluable (**InputPlug.isEvaluable()**, Alg. 17), if it is set to a constant value (**isConstant()**, Alg. 16), or it is connected and it does not depend on any unconnected (or not constant) input plugs (which are accordingly titled upstream unconnected plugs). If all the input plugs of an action operator are evaluable, the operator itself is also evaluable (**Action.isEvaluable()**, Alg. 18). A behaviour frame is considered evaluable, if its activation plug, *p_{trigger}*, is evaluable (**BehaviourFrame.isEvaluable()**, Alg. 19).

The *isEvaluable()*-functions ensure that the respective plugs and operators have the potential to receive non-*nil* values, either because they are fully upstream connected or because they provide constant, non-*nil* values. This does not, however, rule out the possibility than an evaluable operator may receive an invalid *nil* value during the actual evaluation process. Only if it receives non-*nil* information, we say the plug or the operator is triggered (**Plug.isTriggered()**, **Action.isTriggered()**, and **BehaviourFrame.isTriggered()**, Algs. 20 to 22).

Output: boolean
return self.isConnected() AND
self.owner.owner.containsOperator(p.getConnectedSelection());
Algorithm 15: *InputPlug.isInternallyConnected()*

Output: boolean
return self.isConnected() == false AND self.value != nil;
Algorithm 16: *InputPlug.isConstant()*

Output: boolean
return self.isConstant() OR (self.isConnected() AND
self.upstreamUnconnectedPlugs.size() == 0);
Algorithm 17: *InputPlug.isEvaluable()*

Output: boolean
for $i := 0, i < self.inputPlugs.size(), i++$ **do**
| p := self.inputPlugs.get(i);
| **if** p.isEvaluable() == false **then return** false;
end
return true;
Algorithm 18: *Action.isEvaluable()*

Output: boolean
return self.p_trigger.isEvaluable();
Algorithm 19: *BehaviourFrame.isEvaluable()*

An evaluable input plug is triggered, if it returns a constant value or if it retrieves a non-*nil* value from its connected output plug. In order to test whether a valid value is retrieved from a connected output plug, the method **InputPlug.computeUpstreamSelections()** (Alg. 42) needs to be called, which computes the list of upstream connected selections in the given order and updates the value of the connected output plug. This method is presented as part of the simulation routines of Section 5. Afterward, the input plug retrieves the updated value from the output plug. If it is non-*nil*, the input plug is considered “triggered”. The method **InputPlug.orderUpstream Selections()** (Alg. 23) prepares an input plug’s upstreamSelections list for computation. It is detailed in the following paragraphs on state updating routines.

Action.isTriggered() first tests whether all the input plugs of an action yield non-*nil* values. If it is a level-0 action, i.e. if it does not contain further internal actions, it is indeed triggered. If, however, it contains further internal actions and possibly internal selection operators, the higher-level action as a whole gets triggered only if all its internally wrapped actions are triggered. Hence, if only one contained internal action is not triggered, the high-level operator is not triggered either. The method recursively works through the action hierarchy.

```

Output: boolean
if self.isConstant() then return true;
self.computeUpstreamSelections();
self.value := self.connection.outputPlug.value;
return self.value != nil;

```

Algorithm 20: *InputPlug.isTriggered()*

```

Output: boolean
for i := 0, i < self.inputPlugs.size(), i ++ do
  | p := self.inputPlugs.get(i);
  | if p.isTriggered() == false then return false;
end
if self.level > 0 then
  | for i := 0, i < self.internalActions.size(), i ++ do
  | | a := self.internalActions.get(i);
  | | if a.isTriggered() == false then return false;
  | end
end
return true;

```

Algorithm 21: *Action.isTriggered()*

Output: boolean
return self.*p_trigger.isTriggered*();
Algorithm 22: *BehaviourFrame.isTriggered()*

3.3 State Updating Routines

The method *InputPlug.orderUpstreamSelections()* (Alg. 23) retrieves and orders the internal selection operators that are connected to an input plug. It also aggregates upstream plugs that are neither constant nor connected. Each internally upstream-connected Selection instance is added to the list of upstreamSelections and is marked with a step-wise increased index value, starting from 0. Its input plugs are added to the list of plugs to follow up on. During the upstream traversal of connected selections, the selections are assigned index values that inversely reflect their order of access—selections with high indices need to be computed before those with lower indices. The list of upstream selections is sorted accordingly. Input plugs that are neither connected nor constant represent dependencies that cannot be internally resolved—these are added to the list of *upstreamUnconnectedPlugs*⁴.

```

self.upstreamUnconnectedPlugs := {};
self.upstreamSelections := {};
index := 0;
inputPlugs := {self};
for i := 0, i < inputPlugs.size(), i ++ do
    pin := inputPlugs.get(i);
    if pin.isInternallyConnected() then
        connectedSelection := pin.getConnectedSelection();
        connectedSelection.index := index ++;
        inputPlugs.append(connectedSelection.inputPlugs);
        self.upstreamSelections.append(connectedSelection);
    else if pin.isConnected() == false AND pin.isConstant() == false then
        self.upstreamUnconnectedPlugs.append(pin);
    end
end
self.upstreamSelections.sort() with descending index;
Algorithm 23: InputPlug.orderUpstreamSelections()

```

Figure 9 depicts the process effected by *InputPlug.orderUpstreamSelections()* working on *B*'s upper input plug (double-lined, red circle). Afterward, its associated list *upstreamUnconnectedPlugs* will contain *A*'s upper input plug.

⁴ External connections, i.e. those that connect to non-internal operators, need not be considered in the *orderUpstreamSelections()* method since (1) external operators will be executed at a higher algorithmic level and (2) any upstream unconnected plugs are connected through exposed plugs of the high-level operator anyways.

The list *upstreamSelections* will contain *C* and *A*, and in this order, since *C* needs to compute first. In Figure 10 the upstream selections of the second input plug (marked with a double red line) of an action operator are computed. Its list *upstreamConnectedPlugs* will contain the blue coloured input plugs of operators *A'* and *D*, its *upstreamSelections* list will be $\{D, G, A', F\}$. During the computation of the upstream selections, the index value of *D* was set to 3 first, then overwritten by the value 4 since two independent downstream selection operators accessed it. Overwriting indices ensures that an operator is computed as soon as necessary.

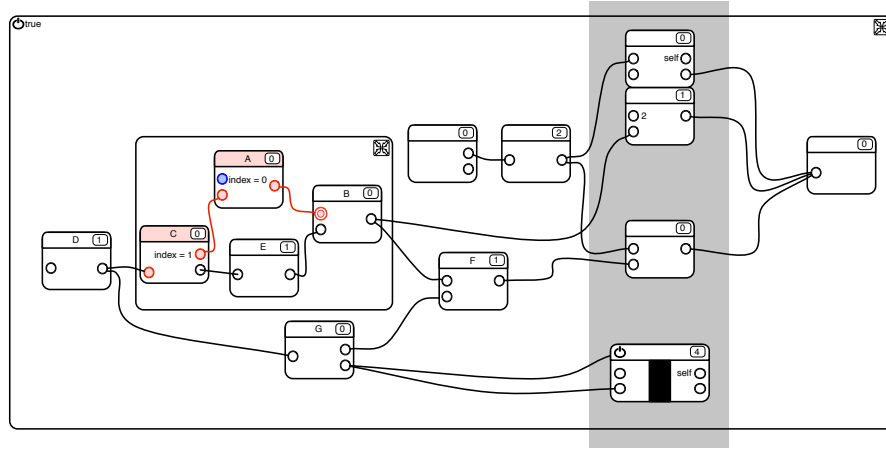


Fig. 9. Starting from the upper one of *B*'s input plugs (the double-lined red circle), the upstream selection operators *A* and *C* are determined and ordered for execution.

4 Routines of the Modelling Phase

The routines of the modelling phase can be roughly divided into three categories: Adding and removing operators, adding and removing connections, and creating high-level operators.

4.1 Adding/Removing Operators

During the modelling phase, selections, actions and behaviour frames are combined to shape the behaviour of an agent. A single behaviour frame, also referred to as the “root” behaviour frame, serves as a workspace to host and configure SwarmScript operators. Predefined, configurable SwarmScript operators may be provided by an operator library and be added to the behaviour frame or to its lower-level operators.

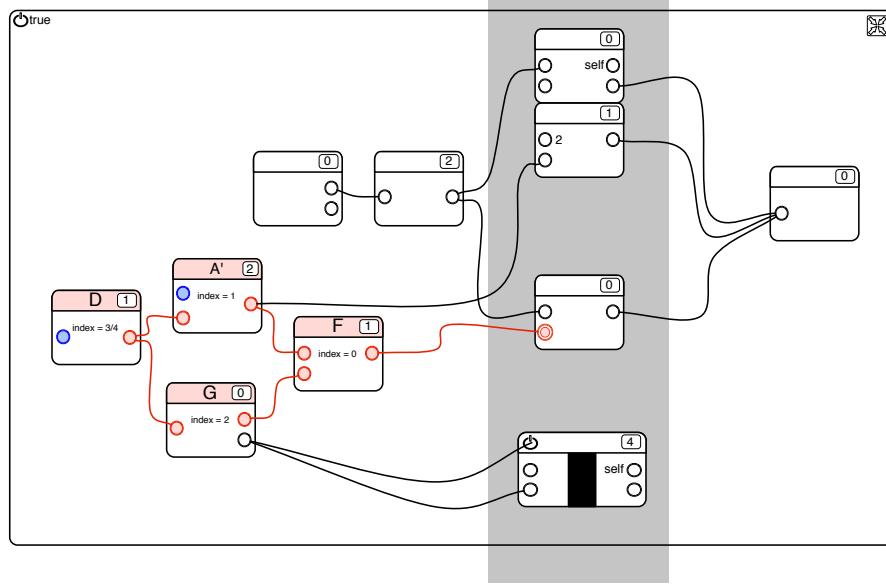


Fig. 10. Starting from an action operator's input plug (the double-lined red circle), the upstream selection operators F , A' , G and D are determined and stored as the ordered list $\{D, G, A', F\}$.

If possible, **Operator.addTo()** (Alg. 24) adds an operator to another one. Behaviour frames can host any kind of operators, but they cannot be added to actions, and neither actions nor behaviour frames can be added to selections. Once a new operator has been added, its owner is updated in order to ensure that it exposes the proper plugs. The implementations of the **Operator.update()** method are presented in the paragraphs about creating high-level operators, starting with Section 4.3.

The transfer of an operator to another enclosing operator is realized by the **Operator.transferTo** (Alg. 25) method. **Operator.delete()** (Alg. 27) completely removes an operator, first deleting the connections of all its plugs, then detaching itself from its owner, and finally removing itself.

4.2 Adding/Removing Connections

Programmatic relationships are created through setting up connections from output to input plugs. Thus, selection operators are joined into chains that are eventually hooked up to the activation plugs $p_{trigger}$ of behaviour frames or to the source and target plugs of action operators. A **Connection** instance references an output and an input plug and is itself added to the connection list of its input and referenced as the connection of an output plug. When connecting an output plug or input plug these dependencies are updated accordingly

```

Input: to: Operator
if class of self is Selection then
  | to.internalSelections.append(self);
else if class of self is Action AND class of to is not Selection then
  | to.internalActions.append(self);
else if class of self is BehaviourFrame AND class of to is BehaviourFrame then
  | to.internalBehaviourFrames.append(self);
end
self.owner := to;
self.owner.update();

```

Algorithm 24: *Operator.addTo()*

```

Input: to: Operator
if class of self is Selection then
  | o.owner.internalSelections.remove(self);
  | to.internalSelections.append(self);
else if class of self is Action AND class of to is not Selection then
  | o.owner.internalActions.remove(self);
  | to.internalActions.append(self);
else if class of self is BehaviourFrame AND class of to is BehaviourFrame then
  | o.owner.internalBehaviourFrames.remove(self);
  | to.internalBehaviourFrames.append(self);
end
self.owner.update();
self.owner := to;
self.owner.update();

```

Algorithm 25: *Operator.transferTo()*

```

if class of self is Selection then
  | foreach plug p ∈ self.outputPlugs do
  | | p.disconnect();
  | end
end
foreach plug p ∈ self.inputPlugs do
  | if p.isConnected() then p.disconnect();
end

```

Algorithm 26: *Operator.disconnect()*

(*OutputPlug.connectTo()*, *InputPlug.connectTo()*, Algs. 29 and 28). If an output plug is connected to an already connected input plug, that input plug is disconnected first.

When a connection is created, the enclosing host operator needs to be updated, as its internal topology has changed. In case, the connection is part of the internal topological definition of a selection operator, the execution sequence of internal operators has to be re-calculated. If the connection is encapsulated by an action operator or inside a behaviour frame, it plays a role in connecting a

```

self.disconnect();
if class of self is Selection then
| self.owner.internalSelections.remove(self);
else if class of self is Action then
| self.owner.internalActions.remove(self);
else if class of self is BehaviourFrame then
| self.owner.internalBehaviourFrames.remove(self);
end
self.owner.update();
delete(self);

```

Algorithm 27: *Operator.delete()*

selection operator or a chain of selection operators to one or more actions. The affected action plugs need to update their dependencies accordingly, as well. No matter which wrapping operator hosted the connection, it might have to update its exposed plugs.

When disconnecting an input plug (*InputPlug.disconnect()*, Alg. 30), the reference to the corresponding Connection instance needs to be removed from the connected output plug. Then, the Connection instance is deleted, also setting the input plug’s reference to *nil*.

Finally, all the affected non-Selection plugs need to update their upstream dependencies, since they have changed. And the associated owner’s host needs to get updated as its exposed plugs might have changed as well. Disconnecting an output plug, requires the removal of all its connections (*OutputPlug.disconnect()*, Algs. 32), including updating the owners of the connected operators as well as updating any affected input plugs that belong to actions and behaviour frames⁵.

```

Input: to: OutputPlug
if self.isConnected() then self.disconnect();
self.connection := new Connection(to, self);
to.connections.append(self.connection);
foreach Plug p ∈ self.getAffectedInputPlugs() do
| p.orderUpstreamSelections();
end
self.owner.update();

```

Algorithm 28: *InputPlug.connectTo()*

4.3 Creating High-Level Operators

In order to provide powerful programmatic building blocks to the modeller, we present routines that wrap SwarmScript “circuitry”, including operators and

⁵ There is no need to update anything inside of the owners of the identified non-Selection plugs, since the non-exposed plugs do not have any external dependencies.

Input: to: InputPlug
to.connectTo(self);

Algorithm 29: *OutputPlug.connectTo()*

```
self.connection.outputPlug.connections.remove(self.connection);
delete(self.connection);
self.connection := nil;
foreach Plug  $p \in self.getAffectedInputPlugs()$  do
  |  $p.orderUpstreamSelections()$ ;
end
```

Algorithm 30: *InputPlug.disconnect()*

```
self.disconnect();
self.owner.owner.update();
```

Algorithm 31: *InputPlug.cleanDisconnect()*

```
foreach InputPlug  $p_{in} \in self.getConnectedInputPlugs()$  do
  |  $p_{in}.disconnect()$ ;
end
```

Algorithm 32: *OutputPlug.disconnect()*

```
self.disconnect();
self.owner.owner.update();
```

Algorithm 33: *OutputPlug.cleanDisconnect()*

connections, into higher-level operators. The plugs exposed by a high-level operator reference the plugs of the internal selections, actions and behaviour frames. Following this approach, the connection dependencies of newly wrapped operators and their contexts are automatically maintained.

wrapOperators() Alg. 34 calls **wrapSelection()** (Alg. 35), if the selected circuitry is limited to selection operators. One action in combination with some associated selections are wrapped by means of the function **wrapAction()** (Alg. 37). In case multiple independent actions or one or more behaviour frames are wrapped, a new behaviour frame needs to be created relying on the function **wrapBehaviourFrame()** (Alg. 39).

When multiple independent actions are made dependent, i.e. when they are combined into a sequences of dependent actions, the method *wrapAction()* is called directly, with the list of further on dependent actions as a parameter. Due to the additional semantics of rendering actions dependent, this case is not covered by the general *wrapOperators* method.

```

Input: operators: list(Operator)
if operators.size() == 0 then return;
actions := {};
selections := {};
behaviourFrames := {};
foreach Operator o ∈ operators do
    if class of o is Selection then selections.append(o);
    if class of o is Action then actions.append(o);
    if class of o is BehaviourFrame then behaviourFrames.append(o);
end
if behaviourFrames.size() > 0 OR actions.size() > 1 then
    | wrapBehaviourFrame(selections, actions, behaviourFrames);
else if actions.size() == 1 then
    | wrapAction(selections, actions.first());
else
    | wrapSelection(selections);
end

```

Algorithm 34: *wrapOperators()*

wrapSelection() As detailed in Alg. 35, an arbitrarily chosen selection operator (the first of an unordered list) determines the owner of the new high-level selection operator, and whether it will be of type source or target. An according new *Selection* instance is created and added to the behaviour, whereas all the chosen operators are transferred from their current position of the operator hierarchy into the newly created selection operator, or wrapper. Internal selection operators that do not initially share the type of the new high-level selection are disconnected before being transferred in order to avoid semantic conflicts in the model.

The internal dependencies and ownerships relationships of the new wrapper are resolved by means of the **Selection.update()** method (Alg. 36). It computes the exposed input and output plugs of the new selection wrapper and sets the owner property of the exposed plugs to reference the new high-level selection operator. The level of recursion of the wrapper is determined by the maximal level of any of the internal operators plus one. In order to ensure that its output plugs can provide values once the operator is computed, the internal dependencies need to be resolved. In particular, the upstream selection chains are computed for the internal input plugs that are owned by those internal operators that deliver the output of the high-level selection operator.

Figure 11 shows these mechanisms in detail. The input plugs coloured in blue are the ones that feed into internal operators that will provide the output plugs of the wrapper. The upstream selections of these blue plugs need to be determined. Any input and output plugs that are not internally connected and not constant are coloured in red (connections to the outside of the operator are indicated by red lines). These plugs are adopted and exposed by the new higher-level operator.


```

Input: selections: list(Selection)
owner := selections.first().owner;
type := selections.first().type;
newSelection := new Selection(owner, 0, {}, {}, type, {}, false, 0);
owner.internalSelections.append(newSelection);
foreach Selection s ∈ selections do
    | if s.type ≠ type then s.disconnect();
    |   s.transferTo(newSelection);
end
newSelection.update();

```

Algorithm 35: wrapSelection()

```

self.level := 0;
self.inputPlugs := {};
self.outputPlugs := {};
foreach Selection s ∈ self.internalSelections do
    | self.level := max(self.level, s.level);
    | s.owner := self;
    | outputOperator := false;
    | foreach OutputPlug p ∈ s.outputPlugs do
    |   | if p.isInternallyConnected() == false then
    |     | p.owner := self;
    |     | self.outputPlugs.append(p);
    |     | outputOperator := true;
    |   end
    | end
    | foreach InputPlug p ∈ s.inputPlugs do
    |   | if outputOperator == true then p.orderUpstreamSelections();
    |   | if p.isConstant() == false AND p.isInternallyConnected == false
    |   |   then
    |   |     | p.owner := self;
    |   |     | self.inputPlugs.append(p);
    |   |   end
    |   end
    | end
end
self.level += 1;
self.owner.update();

```

Algorithm 36: Selection.update()

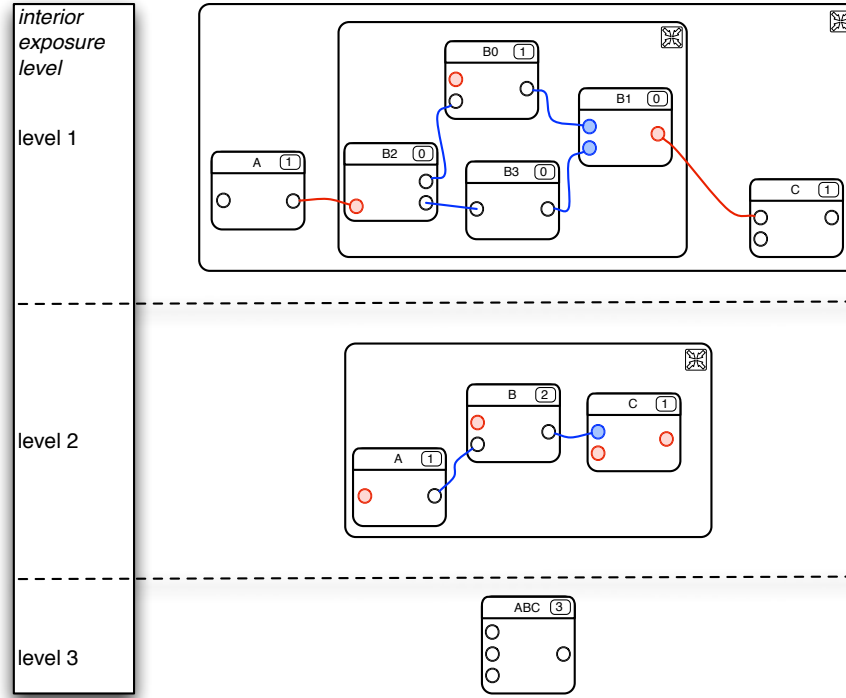


Fig. 11. First the four interior selection operators $B0...B3$ are wrapped into level-2 operator B . Next, A , B , and C are wrapped into the operator ABC . Red plugs at one level become the exposed plugs at the next higher level. Blue plugs denote those internal input plugs whose upstream dependencies need to be pre-calculated before the wrapper can be computed.

wrapAction() Alg. 37 wraps a combination of selection operators and an action or an action sequence into a higher-level action operator. It is important to understand that feeding multiple actions into the *wrapAction()* method assumes that they are dependent—they will be wrapped into one new high-level action, and thus either none of them or all of them will be executed (in the given order).

An action operator does not expose any output plugs since it does not provide information. It rather creates, deletes, or changes information. Therefore, unlike *Selection.update()*, **Action.update()** (Alg. 38) does not consider the output plugs of any wrapped selections. It solely exposes a wrapped selection's input plugs, if they are neither constant nor internally connected. Depending on the Selection instance's type (source or target), any adopted input plugs become source or target plugs of the wrapper, respectively.

If the plug of an internal action is internally connected, its internal dependencies are computed by means of the method *orderUpstreamSelections()*. If it is externally connected or not connected at all, it will get exposed by the wrapper.

Updating a behaviour frame is fairly similar to wrapping an action, which is why *Action.update()* covers updating behaviour frames as well. Wrapped behaviour frames are treated exactly as wrapped actions, except for the additional consideration of the special activation plug *ptrigger* of a behaviour frame. Since it algorithmically works like a regular source input plug, it is added to the list of input plugs and processed, i.e. ignored or adopted, like the other ones.

```

Input: selections: list(Selection), actions: (Action)
owner := actions.first().owner;
newAction := new Action(owner, 0, {}, {}, {});
foreach Selection s ∈ selections do
  | s.transferTo(newAction);
end
foreach Action a ∈ actions do
  | a.transferTo(newAction);
end
newAction.update();

```

Algorithm 37: *wrapAction()*

wrapBehaviourFrame() Wrapping operators into a BehaviourFrame instance is very similar to creating a higher level Action instance as described in the previous paragraph. However, there are two subtle, but important differences regarding the semantics of behaviour frames.

1. A behaviour frame can host multiple independent actions. As a consequence, a behaviour frame does not follow the same algorithmic rules as an action. As the source and target plugs exposed by a behaviour frame may stem from independent internal actions, insisting on valid input values before considering all the wrapped actions for execution would be detrimental.
2. It is the purpose of behaviour frames to allow the modeller to manage groupings of arbitrary behavioural “circuitry”. For this reason, they provide an on/off switch for the consideration of the wrapped behavioural code (Section 2.8). A behaviour frame exposes an according special activation plug *ptrigger* that is set to a valid value—*true* in the given case—by default which activates the behaviour frame. Like a source plug of an action, *ptrigger* can be connected to an arbitrary selection, and an input value of *nil* will prevent the consideration of the behaviour frame’s contents.

Accordingly, Algorithm 39 instantiates the Plug instance *ptrigger* of type source, with its value set to *true*. Then the function passes *ptrigger* to a newly created

```

self.level := 0;
self.inputPlugs := {};
foreach Selection s  $\in$  self.internalSelections do
    self.level := max(self.level, s.level);
    s.owner := self;
    for i := 0, i < s.inputPlugs.size(), i ++ do
        p := s.inputPlugs.get(i);
        if p.isConstant() OR p.isInternallyConnected() then continue;
        p.type := s.type;
        p.owner := self;
        self.inputPlugs.append(p);
    end
end
informedOperators := self.internalActions.copy();
if class of self is BehaviourFrame then
    informedOperators.append(self.internalBehaviourFrames);
end
foreach Operator o  $\in$  informedOperators do
    self.level := max(self.level, o.level);
    o.owner := self;
    internalInputPlugs := {};
    if class of o is BehaviourFrame then internalInputPlugs.append(o.ptrigger);
    internalInputPlugs.append(o.inputPlugs);
    for i := 0, i < internalInputPlugs.size(), i ++ do
        p := internalInputPlugs.get(i);
        if p.isConstant() then continue;
        if p.isInternallyConnected() then
            p.orderUpstreamSelections();
        else
            p.owner := self;
            self.inputPlugs.append(p);
        end
    end
end
self.level++ = 1;
self.owner.update();

```

Algorithm 38: *Action.update()*




Fig. 12. The combination of interior action and selection operators is wrapped into a level-3 action which exposes five plugs (coloured red at level-2). The blue plug and blue connection on level 2 denote the only internal dependencies of the level-3 action. The process is repeated to forge a level-4 action.

behaviour frame. The remainder of the function steps through the chosen operators, transfers them to the corresponding containers, and finally updates the new BehaviourFrame instance in order to resolve its internal dependencies and to compute its exposed input plugs. The update of a BehaviourFrame instance is performed by the *Action.update()* method discussed above.

5 Simulation Routines

Before entering the simulation phase, the swarm script behaviour can be optimized by pruning away actions and behaviour frames that are not evaluable, and thus would not impact any simulation processes.

```

Input: selections list⟨Selection⟩, actions: ⟨Action⟩, behaviourFrames:
        ⟨BehaviourFrame⟩
if actions.size() > 0 then
    | owner := actions.first().owner;
else
    | owner := behaviourFrames.first().owner;
end
ptrigger := new Plug(nil, true, source, nil, {}, {});
newBehaviourFrame := new BehaviourFrame(owner, 0, {}, {}, {}, ptrigger, {});
ptrigger.owner := newBehaviourFrame;
foreach Selection s ∈ selections do
    | s.transferTo(newBehaviourFrame);
end
foreach Action a ∈ actions do
    | a.transferTo(newBehaviourFrame);
end
foreach BehaviourFrame bf ∈ behaviourFrames do
    | bf.transferTo(newBehaviourFrame);
end
newBehaviourFrame.update();

```

Algorithm 39: *wrapBehaviourFrame()*

Each simulation step can further be divided into several sub-steps: (1) Computing the chains of selections that feed into the actions and behaviour frames, (2) identifying triggered actions and behaviour frames, and (3) the actualization of the triggered actions.

5.1 Initialization

During the simulation phase, the agents' interactions are computed. Since the modelling routines update the operators' dependencies, the simulation process does not require further initialization. However, actions and behaviour frames that are not evaluable can be excluded from the set of operators considered for execution upfront. The BehaviourFrame method **BehaviourFrame.updateEvaluableOperators()** (Alg. 40) filters the operators contained in a behaviour frame accordingly, adding only evaluable operators to its *evaluableActions* and *evaluableBehaviourFrames* containers, respectively. Internally hosted behaviour frames are filtered recursively, resulting in a complete optimization of the Swarm-Script code.

5.2 Compute Selections

At each simulation step, the filtered operator hierarchy is searched for actions and behaviour frames that are successfully triggered. In order to get triggered, the input plugs of actions need to receive valid values (non *nil*), as do the activation plugs *ptrigger* of behaviour frames.

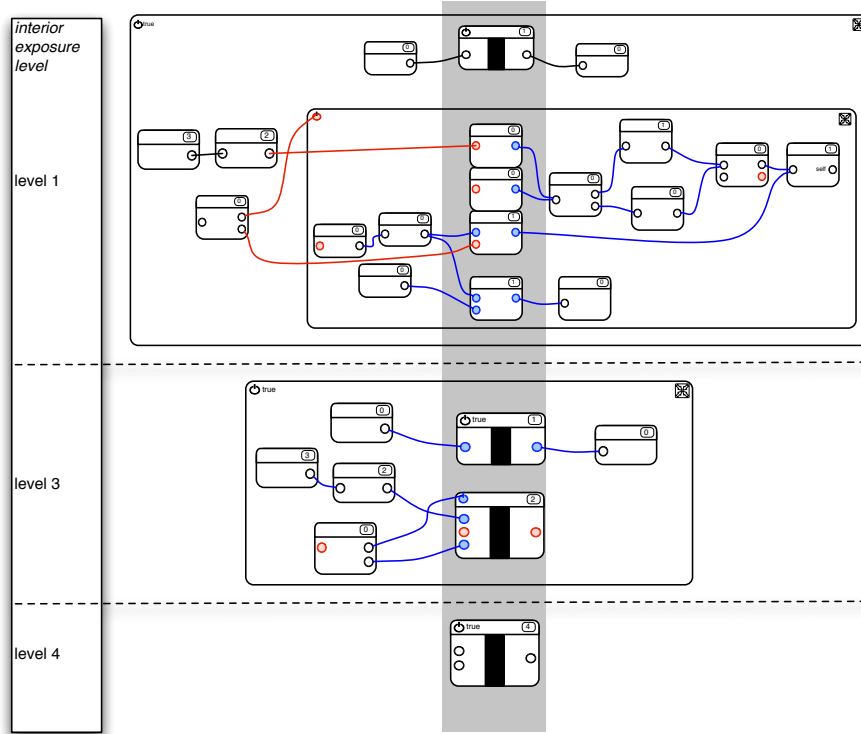


Fig. 13. The inner operators are wrapped into a level-2 behaviour frame. In addition to three source plugs and one target plug, the wrapper exposes its activation plug $p_{trigger}$. External connections are coloured in red, blue connections denote the computation of upstream dependencies of the blue coloured input plugs.

The computation of a selection (**Selection.compute()**, Alg. 41) is performed recursively on its hierarchical structure. At each level, the method **Selection.compute()** (Alg. 41) steps through the operator's internally ordered upstream selections that were calculated during the modelling phase by means of the method *orderUpstreamSelections()*. On each of those input plugs the method **InputPlug.computeUpstreamSelections()** (Alg. 42) is performed ensuring that all the necessary, upstream linked selections are computed in order and provide the selection operator's results. A non-hierarchically defined selection operator (level-0), overrides the *Selection.compute()* method to store values in its output plugs, considering the values of possible input plugs.

5.3 Identify Triggered Actions

At this point of the simulation step, all the necessary selection values are provided and it can be decided which behaviour frames and actions are triggered.

```

self.evaluableActions := {};
self.evaluableBehaviourFrames := {};
for  $i := 0, i < self.internalActions.size(), i++$  do
  |  $a := self.internalActions.get(i);$ 
  | if  $a.isEvaluable()$  then  $self.evaluableActions.append(a);$ 
end
for  $i := 0, i < self.internalBehaviourFrames.size(), i++$  do
  |  $bf := self.internalBehaviourFrames.get(i);$ 
  | if  $bf.p_{trigger}.isEvaluable()$  then
    |  $self.evaluableBehaviourFrames.append(bf);$ 
    |  $bf.updateEvaluableOperators();$ 
  | end
end

```

Algorithm 40: *BehaviourFrame.updateEvaluableOperators()*

```

outputSelections := {};
foreach  $Plug\ p \in self.outputPlugs$  do
  | if  $outputSelections.hasElement(p.owner) == false$  then
    |  $outputSelections.append(p.owner);$ 
end
foreach  $Selection\ s \in outputSelections$  do
  | for  $i := 0, i < s.inputPlugs.size(), i++$  do
    |  $plug := s.inputPlugs.get(i);$ 
    |  $plug.computeUpstreamSelections();$ 
  | end
end
self.computed := true;

```

Algorithm 41: *Selection.compute()*

```

index := self.upstreamSelections.size() - 1;
while  $index \geq 0$  AND  $self.upstreamSelections.get(index).computed == false$ 
do
  |  $index--;$ 
end
 $index++;$ 
for  $i := index, i < self.upstreamSelections.size(), i++$  do
  |  $self.upstreamSelections.get(i).compute();$ 
end

```

Algorithm 42: *InputPlug.computeUpstreamSelections()*

As a result, the hierarchical operator structure is reduced to a list of triggered actions by means of the function **BehaviourFrame.getTriggeredActions()** (Alg. 43). It identifies and collects the triggered actions recursively contained in any considered behaviour frames.

```

Output:  $\langle Action \rangle$ 
triggeredActions := {};
for  $i := 0, i < self.evaluableActions.size(), i++$  do
    |  $a := self.evaluableActions.get(i);$ 
    | if  $a.isTriggered()$  then triggeredActions.append(a);
end
for  $i := 0, i < self.evaluableBehaviourFrames.size(), i++$  do
    |  $bf := self.evaluableBehaviourFrames.get(i);$ 
    | if  $bf.p_{trigger}.isTriggered()$  then
        | triggeredActions.append(bf.getTriggeredActions());
    end
return triggeredActions;
Algorithm 43: BehaviourFrame.getTriggeredActions()

```

5.4 Compute and Actualize Triggered Actions

Next, the triggered operators are executed. High-level actions recursively iterate through any internal actions. Level-0 actions, similar to level-0 selections, need to override the **Action.compute()** and the **Action.actualize()** methods (Algs. 44 and 45). First, these methods process the information the operator gathers from the source inputs and second, they manipulate the target inputs. Accordingly, **Action.compute()** computes the necessary information, and **Action.actualize()** actualizes the changes of the action.

Asynchronous actualization (Alg. 46) iterates through the triggered actions one after another, computes their required parameters based on the provided selection values and actualizes their changes. This approach does not consider the impact of step-wise state changes on subsequent actions.

Using synchronous actualization (Alg. 47), the calculation of the actions' parameters is concluded during the *Action.compute()* method, before any changes are introduced into the system by means of the *Action.actualize()* method. Furthermore, in order to effectively simulate concurrency, the actualization of state changes needs to be ordered based on a given priority policy.

5.5 Running the Simulation

Before running the simulation, all the given behaviour frames that store all the agents' behaviours are filtered for evaluable operators by means of the

```

if self.level > 0 then
  for i := 0, i < self.internalActions.size(), i ++ do
    a := self.internalActions.get(i);
    a.compute();
  end
else
  self.compute();
end

```

Algorithm 44: *Action.compute()*

```

if self.level > 0 then
  for i := 0, i < self.internalActions.size(), i ++ do
    a := self.internalActions.get(i);
    a.actualize();
  end
else
  self.actualize();
end

```

Algorithm 45: *Action.actualize()*

```

Input: actions: list(Action)
for i := 0, i < actions.size(), i ++ do
  actions.get(i).compute();
  actions.get(i).actualize();
end

```

Algorithm 46: *asynchronousActualize()*

```

Input: actions: list(Action)
sort actions in accordance with order policy  $\mathcal{P}$ ;
for i := 0, i < actions.size(), i ++ do
  actions.get(i).compute();
end
for i := 0, i < actions.size(), i ++ do
  actions.get(i).actualize();
end

```

Algorithm 47: *synchronousActualize()*

updateEvaluableOperators() method. Then, while the simulation is not terminated, all those optimized behaviours are considered for execution at each step (Alg. 48).

```

Input: behaviours: list(BehaviourFrame)
foreach BehaviourFrame bf in behaviours do
  | bf.updateEvaluableOperators();
end
while not terminated do
  | stepSimulation(behaviours);
end

```

Algorithm 48: *runSimulation()*

The procedure **stepSimulation()** (Alg. 49) steps through a list of behaviour frames, computes the triggered actions and actualizes them. As explained in the preceding paragraphs, the selections linked to the actions' plugs need to be computed in order to verify whether an action is actually triggered or not. Therefore, all the three steps—(1) selection computation, (2) triggered action identification, and (3) triggered action actualization—are performed by the *stepSimulation()* procedure.

```

Input: behaviours: list(BehaviourFrame)
forall the Selection instances s do
  | s.computed := false;
end
actionCache := {};
foreach BehaviourFrame bf in behaviours do
  | actionCache.append(bf.getTriggeredActions());
end
actualize(actionCache); (*synchronous or asynchronous*)

```

Algorithm 49: *stepSimulation()*