

Compositional Reasoning about Active Objects with Shared Futures^{*}

Crystal Chang Din¹ and Olaf Owe²

¹ Department of Computer Science, Technische Universität Darmstadt, Germany
crystalcd@cs.tu-darmstadt.de

² Department of Informatics, University of Oslo, Norway
olaf@ifi.uio.no

Abstract. Distributed and concurrent object-oriented systems are difficult to analyze due to the complexity of their concurrency, communication, and synchronization mechanisms. The *future mechanism* extends the traditional method call communication model by facilitating sharing of references to futures. By assigning method call result values to futures, third party objects may pick up these values. This may reduce the time spent waiting for replies in a distributed environment. However, futures add a level of complexity to program analysis, as the program semantics becomes more involved.

This paper presents a model for asynchronously communicating objects, where return values from method calls are handled by futures. The model facilitates invariant specifications over the locally visible communication history of each object. Compositional reasoning is supported and proved sound, as each object may be specified and verified independently of its environment. A kernel object-oriented language with futures inspired by the ABS modeling language is considered. A compositional proof system for this language is presented, formulated within dynamic logic.

1 Introduction

Distributed systems play an essential role in society today. However, quality assurance of distributed systems is non-trivial since they depend on unpredictable factors, such as different processing speeds of independent components. Therefore, it is highly challenging to test such distributed systems after deployment under different relevant conditions. These challenges motivates frameworks combining precise modeling and analysis with suitable tool support. In particular, *compositional verification systems* allow the different components to be analyzed independently from their surrounding components.

Object orientation is the leading framework for concurrent and distributed

^{*} Supported by the EU project FP7-610582 *Envisage: Engineering Virtualized Services* (<http://www.envisage-project.eu>) and FP7-ICT-2013-X *UpScale: From Inherent Concurrency to Massive Parallelism through Type-based Optimizations* (<http://www.upscale-project.eu>).

systems, recommended by the RM-ODP [26]. However, method-based communication between concurrent units may cause busy-waiting, as in the case of remote and synchronous method invocation, e.g., Java RMI [3]. Concurrent objects communicating by *asynchronous method calls* have been proposed as a promising framework to combine object-orientation and distribution in a natural manner. Each concurrent object encapsulates its own state and processor, and internal interference is avoided as at most one process is executing on an object at a time. Asynchronous method calls allow the caller to continue with its own activity without blocking while waiting for the reply, and a method call leads to a new process on the called object. The notion of *futures* [6, 23, 30, 35] improves this setting by providing a decoupling of the process invoking a method and the process reading the returned value. By sharing *future identities*, the caller enables other objects to wait for method results. However, futures complicate program analysis since programs become more involved compared to semantics with traditional method calls, and in particular local reasoning is a challenge.

The execution of a distributed system can be represented by its *communication history* or *trace*; i.e., the sequence of observable communication events between system components [8, 25]. At any point in time the communication history abstractly captures the system state [11, 12]. In fact, traces are used in semantics for full abstraction results (e.g., [1, 27]). The *local history* of an object reflects the communication visible to that object, i.e., between the object and its surroundings. A system may be specified by the finite initial segments of its communication histories, and a *history invariant* is a predicate which holds for all finite sequences in the set of possible histories, expressing safety properties [5].

In this work we consider a kernel object-oriented language, where futures are used to manage return values of method calls. Objects are concurrent and communicate asynchronously. We formalize object communication by a four event operational semantics, capturing shared futures, where each event is visible to only one object. Consequently, the local histories of two different objects share no common events, and history invariants can be established independently for each object. We present a *dynamic logic* proof system for class verification, facilitating independent reasoning about each class. A verified *class invariant* can be instantiated to each object of that class, resulting in an invariant over the local history of the object. Modularity is achieved as the independently derived history invariants can be composed to form *global* system specifications. Global history consistency is captured by a notion of history *well-formedness*. The formalization of object communication extends previous work [18] which considered concurrent objects and asynchronous communication, but without futures.

Paper overview. Sect. 2 presents a core language with shared futures. The communication model is presented in Sect. 3, and Sect. 4 defines the operational semantics. Sect. 5 presents the compositional reasoning system, and Sect. 6 contains related work and concludes the paper.

$Cl ::= \mathbf{class} C([T cp]^*) \{[T w [:= e]^?]^* s M^*\}$	class definition
$M ::= T m([T x]^*) \{\mathbf{var} [T x]^? s; \mathbf{return} e\}$	method definition
$T ::= C \mid \mathbf{Int} \mid \mathbf{Bool} \mid \mathbf{String} \mid \mathbf{Void} \mid \mathbf{Fut} \langle T \rangle$	types
$v ::= x \mid w$	variables (local or field)
$e ::= \mathbf{null} \mid \mathbf{this} \mid v \mid cp \mid f(\bar{e})$	pure expressions
$s ::= v := e \mid fr := v!m(\bar{e}) \mid v := e?$	statements
$\quad \mid \mathbf{skip} \mid \mathbf{if} e \mathbf{then} s [\mathbf{else} s]^? \mathbf{fi} \mid s; s$	
$\quad \mid \mathbf{while} e \mathbf{do} s \mathbf{od} \mid v := \mathbf{new} C(\bar{e})$	

Fig. 1. Core language syntax, with C class name, cp formal class parameter, m method name, w fields, x method parameter or local variable, and where fr is a future variable. We let $[]^*$ and $[]^?$ denote repeated and optional parts, respectively, and \bar{e} is a (possibly empty) expression list. Expressions e and functions f are side-effect free.

2 A Core Language with Shared Futures

A *future* is a placeholder for the return value of a method call. Each future has a unique identity which is *generated* when a method is invoked. The future is *resolved* upon method termination, by placing the return value of the method in the future. Thus, unlike the traditional method call mechanism, the callee does not send the return value directly back to the caller. However, the caller may keep a *reference* to the future, allowing the caller to *fetch* the future value once resolved. References to futures may be shared between objects, e.g., by passing them as parameters. After achieving a future reference, this means that third party objects may fetch the future value. Thus, the future value may be fetched several times, possibly by different objects. In this manner, shared futures provide an efficient way to distribute method call results to a number of objects.

For the purposes of this paper, we consider a core object-oriented language with futures, presented in Fig 1. It includes basic statements for first class futures, inspired by *ABS* [24]. Class instances are concurrent, encapsulating their own state and processor. Each method invoked on the object leads to a new process, and at most one process is executing on an object at a time. Object communication is *asynchronous*, as there is no explicit transfer of control between the caller and the callee. Methods are organized in classes in a standard manner. A class C takes a list of formal parameters \bar{cp} , defines fields \bar{w} , initialization block s and methods \bar{M} . There is read-only access to the parameters \bar{cp} . A method definition has the form $m(\bar{x})\{\mathbf{var} \bar{y}; s; \mathbf{return} e\}$, ignoring type information, where \bar{x} is the list of parameters, \bar{y} an optional list of *method-local variables*, s is a sequence of statements, and the value of e is returned upon termination.

A future variable fr is declared by $\mathbf{Fut} \langle T \rangle fr$, indicating that fr may refer to futures which will eventually contain values of type T . The call statement $fr := x!m(\bar{e})$ invokes the method m on object x with input values \bar{e} . The identity of the generated future is assigned to fr , and the calling process continues execution without waiting for fr to become resolved. The query statement $v := fr?$ is used to fetch the value of a future. The statement blocks until fr is resolved, and

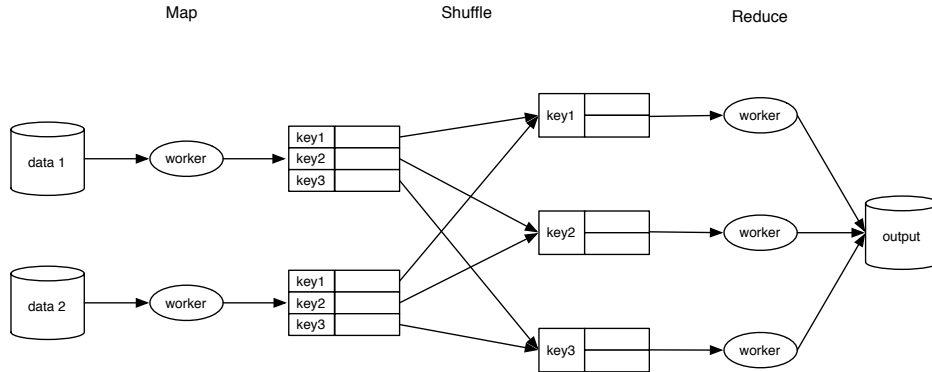


Fig. 2. MapReduce model.

then assigns the value contained in fr to v . The language contains additional statements for assignment, **skip**, conditionals, and sequential composition.

We assume that call and query statements are well-typed. If x refers to an object where m is defined with no input values and return type Int , the following is a well-typed blocking method call: $\text{Fut} \langle \text{Int} \rangle fr; \text{Int } v; fr := x!m(); v := fr?$.

To avoid blocking, *ABS* provides statements for process control, including a statement **await** $fr?$, which releases the current process as long as fr is not yet resolved. This gives rise to more efficient computing with futures. It is possible to add a treatment of process release statements as a straight forward extension of the present work, following the approach of [18]. We here focus on a core language for futures, with a simple semantics, avoiding specialized features such as process control. The core language ignores *ABS* features that are orthogonal to shared futures, including interface encapsulation, inheritance, local synchronous calls, and internal scheduling of processes by means of cooperative multitasking. We refer to the report version of this paper for a treatment of these issues [19].

2.1 The MapReduce Example

In order to illustrate the usage of futures, we consider the problem of counting the number of occurrences of each word in a large collection of documents. We consider the computing model MapReduce in Fig. 2. MapReduce is invented and used heavily by Google for efficient distributed computing over large data sets [17]. It has three major steps: Map, Shuffle and Reduce. The Map phase runs over input data, which might be a database or some files, and output key-value pairs. The input data is split in parts so they can be processed by workers in parallel. The second step is the Shuffle phase, which collates values with the same key together. At last, the Reduce function is called by workers in parallel on the shuffled data distinguished by keys.

We assume two interfaces, *WorkerI* and *MapReduceI*. The interface *WorkerI* is implemented by a class *Worker* shown in Listing 1.1, in which the method *invokeMap* takes a file and emits a list of pairs such that each word in the file is associated with a counting number: ‘1’ in this example. For instance, if the content of the file is ‘I am fine’, the output of *invokeMap* is ‘(I,1),(am,1),(fine,1)’. The method *invokeReduce* in class *Worker* sums together all counts emitted for a particular word. For instance, if the word “am” appears twice, *invokeReduce* takes ‘(am, (1,1))’ and outputs 2.

The interface *MapReduceI* is implemented by class *MapReduce*, shown in Listing 1.2. We here assume generic data types for sets, lists, and pairs, the latter with *fst* and *snd* to extract the first and second element, respectively.

The input to the method *mapReduce* is a list of files each starts with a filename and contains a list of words, i.e. the content of the file. Each file are handled by a worker in parallel. To achieve concurrency, for each file the object of *MapReduce* calls asynchronously the method *invokeMap* on the assigned worker *w*. This is realized by the statement *fMap := w!invokeMap(filename, content)*. The function *insertElement* collects all the futures into a set *fMapResults*. Next is the Shuffle phase. The function *take* randomly extracts an element from a set. The method *mapReduce* waits upon each future, gets the results from each future: *mapResult := fMapResult?*, and collates all the values with the same key, i.e. word, together. For instance, ‘(I,1),(am,1),(who,1),(I,1),(am,1)’ is shuffled to ‘(I,(1,1)),(am,(1,1)),(who,(1))’. In the Reduce phase, each ‘key’ and the corresponding values are handled by a worker in parallel. In the same way as the Map phase for achieving concurrency, the first part of the reduce phase calls asynchronously the method *invokeReduce* on the assigned worker *w*. This is realized by the statement *fReduce := w!invokeReduce(key, values)*. The function *insertElement* collects all the futures into a set *fReduceResults*. At the very last, the method *mapReduce* waits upon each future, gets the results from each

```

class Worker () implements WorkerI {
    List<Pair<String, Int>>
        invokeMap(String filename, List<String> content) {...}

    Int invokeReduce(String key, List<Int> value) {...}
    ...
}

class WorkerPool() implements WorkerPoolI {
    WorkerI getWorker() { // provides idle workers,
                        // or generates new workers if needed. }
}

```

Listing 1.1. Sketch of the classes *Worker* and *WorkerPool*.

```

class MapReduce(WorkerPoolI wp) implements MapReduceI {

List<Pair<String, Int>> mapReduce(
List<Pair<String, List<String>>> files) {

Set<Fut<List<Pair<String, Int>>>> fMapResults := EmptySet;
Set<Pair<String, Fut<Int>>> fReduceResults := EmptySet;
List<Pair<String, Int>> result := Nil;

// Map phase //
while (~isEmpty(files)) do
...
WorkerI w := wp.getWorker();
...
Fut<List<Pair<String, Int>>>
  fMap := w!invokeMap(filename, content);
fMapResults := insertElement(fMapResults, fMap)
od;

// Shuffle phase //
while (~emptySet(fMapResults)) do
Fut<List<Pair<String, Int>>>
  fMapResult := take(fMapResults);
...
List<Pair<String, Int>> mapResult := fMapResult?;
... // collates values with the same key together
od;

// Reduce phase //
while (~emptySet(keys)) do
...
WorkerI w := wp.getWorker();
Fut<Int> fReduce := w!invokeReduce(key, values);
fReduceResults := insertElement(
  fReduceResults, Pair(key, fReduce)) od;
while (~emptySet(fReduceResults)) do
Pair<String, Fut<Int>> reduceResult := take(fReduceResults);
...
String key := fst(reduceResult);
Fut<Int> fValue := snd(reduceResult);
Int value := fValue?;
result := Cons(Pair(key, value), result) od;
return result;
}
}

```

Listing 1.2. The MapReduce class. Here the notation $x := o.m(\bar{e})$ abbreviates $u := o!m(\bar{e}); x := u?$ (for some fresh future u) to de-emphasize trivial usage of futures.

future: $value := fValue?$, and return the number of occurrences of each word in a large collection of files.

Here the future mechanism is exploited to make an efficient implementation, avoiding blocking calls on the workers: The Map phase is not waiting for the workers to do *invokeMap*, and is storing future identities only, thereby allowing many workers to start and work concurrently. Likewise in the loop calling *invokeReduce*, only futures identities are recorded. Blocking is delayed to phases where the future value information is actually needed.

2.1.1 The Intentional Reasoning about the MapReduce Example

The implementation of MapReduce must guarantee the accuracy of the output. Namely, the summation of the occurrences of each word in the collection of documents is correct. However, it is not straight forward to verify this system property. The steps of calculation take place in different components in parallel: the Worker objects execute either the Map phase or the Reduce phase, and the MapReduce object shuffles the data and collects the result from the Worker objects. If we only prove the functional correctness of each class, it is not strong enough to prove this system property. Compositional reasoning is therefore required. We need a formalism to capture the interaction (order) between the components such that we are able to derive the system property from the local reasoning of each components. In the end of this paper, we will present a compositional proof of MapReduce which does provide correct number of occurrences of each word in the collection of documents.

3 Observable Behavior

In this section we describe a communication model for concurrent objects communicating by means of asynchronous message passing and futures. The model is defined in terms of the *observable* communication between objects in the system. We consider how the execution of an object may be described by different *communication events* which reflect the observable interaction between the object and its environment. The observable behavior of a system is described by communication histories over observable events [8, 25].

3.1 Communication Events

Since message passing is asynchronous, we consider separate events for method invocation, reacting upon a method call, resolving a future, and for fetching the value of a future. Each event is observable to only one object, which is the one that *generates* the event. The events generated by a method call cycle is depicted in Fig. 3. The object o calls a method m on object o' with input values \bar{e} and where u denotes the future identity. An invocation message is sent from o to o' when the method is invoked. This is reflected by the *invocation event* $\langle o \rightarrow o', u, m, \bar{e} \rangle$ generated by o . An *invocation reaction event* $\langle o \rightarrow o', u, m, \bar{e} \rangle$ is generated by o' once the method starts execution. When the method terminates,




Fig. 3. A method call cycle: object o calls a method m on object o' with future u . The events on the left-hand side are visible to o , those in the middle are visible to o' , and the ones on the right-hand side are visible to o'' . There is an arbitrary delay between message receiving and reaction.

the object o' generates the *future event* $\langle \leftarrow o', u, m, e \rangle$. This event reflects that u is resolved with return value e . The *fetching event* $\langle o \leftarrow, u, e \rangle$ is generated by o when o fetches the value of the resolved future. Since future identities may be passed to other objects, e.g. o'' , this object may also fetch the future value, reflected by the event $\langle o'' \leftarrow, u, e \rangle$, generated by o'' . The creation of an object o' by an object o is reflected by the event $\langle o \xrightarrow{new} o', C, \bar{e} \rangle$, where o' is the instance of class C and \bar{e} are the actual values for the class parameters. Let type **Mid** include all method names, and let **Data** be the supertype of all values occurring as actual parameters, including future identities **Fid** and object identities **Oid**.

Definition 1. (Events) Let $caller, callee, receiver : Oid$, $future : Fid$, $method : Mid$, $class : Cls$, $args : List[Data]$, and $result : Data$. Communication events Ev include:

- Invocation events $\langle caller \rightarrow callee, future, method, args \rangle$, generated by caller.
- Invocation reaction events $\langle caller \rightarrow callee, future, method, args \rangle$, generated by callee.
- Future events $\langle \leftarrow callee, future, method, result \rangle$, generated by callee.
- Fetching events $\langle receiver \leftarrow, future, result \rangle$, generated by receiver.
- Creation events $\langle caller \xrightarrow{new} callee, class, args \rangle$, generated by caller.

Events may be decomposed by functions. For instance, $_.result : Ev \rightarrow Data$ is well-defined for future and fetching events, e.g., $\langle \leftarrow o', u, m, e \rangle.result = e$.

For a method invocation with future u , the ordering of events depicted in Fig. 3 is described by the following regular expression (using \cdot for sequential composition of events)

$$\langle o \rightarrow o', u, m, \bar{e} \rangle \cdot \langle o \rightarrow o', u, m, \bar{e} \rangle \cdot \langle \leftarrow o', u, m, e \rangle [\cdot \langle _ \leftarrow, u, e \rangle]^*$$

for some fixed o, o', m, \bar{e}, e , and where $_$ denotes an arbitrary value. This implies that the result value may be read several times, each time with the same value, namely that given in the preceding future event.

3.2 Communication Histories

The execution of a system up to present time may be described by its history of observable events, defined as a sequence. A sequence over some type T is constructed by the empty sequence ε and the right append function $_ \cdot _ : \text{Seq}[T] \times T \rightarrow \text{Seq}[T]$ (where “ $_$ ” indicates an argument position). The choice of constructors gives rise to generate inductive function definitions, in the style of [12]. Projection, $_ / _ : \text{Seq}[T] \times \text{Set}[T] \rightarrow \text{Seq}[T]$ is defined inductively by $\varepsilon / s \triangleq \varepsilon$ and $(a \cdot x) / s \triangleq \mathbf{if} \ x \in s \ \mathbf{then} \ (a/s) \cdot x \ \mathbf{else} \ a/s \ \mathbf{fi}$, for $a : \text{Seq}[T]$, $x : T$, and $s : \text{Set}[T]$, restricting a to the elements in s . We use dot notation to extract components from record-like structures, for instance $\langle o \rightarrow o', f, m, \bar{e} \rangle$.*callee* is o' , and also lift the dot notation to sequences. For a sequence h of events, h / \leftarrow is the subsequence of invocation events, and (h / \leftarrow) .*callee* is the sequence of callee elements from these invocation events.

A *communication history* for a set S of objects is defined as a sequence of events generated by the objects in S . We say that a history is *global* if S includes all objects in the system.

Definition 2. (Communication histories) *The communication history h of a system of objects S is a sequence of type $\text{Seq}[Ev]$, such that each event in h is generated by an object in S .*

We observe that the *local history* of a single object o is achieved by restricting S to the single object, i.e., the history contains only elements generated by o . For a history h , we let h/o abbreviate the projection of h to the events generated by o . Since each event is generated by only one object, it follows that the local histories of two different objects are disjoint.

Definition 3. (Local histories) *For a global history h and an object o , the projection h/o is the local history of o .*

4 Operational Semantics

Rewriting logic [31] is a logical and semantic framework in which concurrent and distributed systems can be specified in an object-oriented style. Unbounded data structures and user-defined data types are defined in this framework by means of equational specifications. Rewriting logic extends membership equational logic with *rewrite rules*, so that in a *rewrite theory*, the *dynamic* behavior of a system is specified as a set of rules on top of its *static* part, defined by a set of equations. Informally, a *labeled conditional rewrite rule* is a transition $l : t \longrightarrow t' \ \mathbf{if} \ cond$, where l is a *label*, t and t' are terms over typed variables and function symbols of given arities, and $cond$ is a condition that must hold for the transition to

take place. Rewrite rules are used to specify local transitions in a system, from a state fragment that matches the pattern t , to another state fragment that is an instance of the pattern t' . Rules are selected nondeterministic if there are at least two rule instantiations with left-hand sides matching overlapping fragments of a term. Concurrent rewriting is possible if the fragments are non-overlapping. Furthermore, matching is made modulo the properties of the function symbols that appear in the rewrite rule, like associativity, commutativity, identity (ACI), which introduces further nondeterminism. The Maude tools [9] allow simulation, state exploration, reachability analysis, and LTL model checking of rewriting logic specifications. The state of a concurrent object system is captured by a *configuration*, which is an ACI multiset of *units* such as objects and messages, and other relevant system parts, which in our case includes futures. Concurrency is then supported in the framework by allowing concurrent application of rules when there are non-overlapping matches of left-hand sides. The following context rule, which is implicit in rewriting logic, describes interleaving semantics (letting G, G_1, G_2 denote subconfigurations):

$$\text{context rule } \frac{G_1 \rightarrow G_2}{G G_1 \longrightarrow G G_2}$$

4.1 Operational Rules

For our purpose, a *configuration* is a multiset of (concurrent) objects, classes, messages, futures, as well as a representation of the global history. We use blank-space as the multiset constructor, allowing ACI pattern matching. Objects have the form **object**($Id : o, A$) where o is the unique identity of the object and A is a set of semantic attributes, including

- $Cl : c$ the class c of the object,
- $Pr : s$ the remaining code s of the active process,
- $Lvar : l$ the local state l of the active process, including method parameters and the implicit future identity **destiny**,
- $Flds : a$ the state a of the fields, including class parameters,
- $Cnt : n$ a counter n used to generate future identities,
- $Mtd : m$ the name m of the current method.

Similarly, classes have the form

$$\mathbf{class}(Id : c, Par : z, Flds : a, Init : s, Mtds : q, Cnt : n)$$

where c is the class name, z the class parameters, a the fields, and s the initialization code. The variable q is a multiset of method definitions of the form

$$(m, \bar{p}, l, s)$$

where m is the method name, \bar{p} is the list of parameters, l contains the local variables (including default values), and s is the code. The counter n in the class is used to generate object identities.

Messages have the form of invocation events as described above. And, a future unit is of the form **fut**($Id : u, Val : v$) where u is the future identity and v is its value. The global history is represented by a unit **hist**(h) where h is finite sequence of events (initially empty). Remark that a system configuration contains exactly one history. The history is included to define the interleaving semantics upon which we derive our history-based reasoning formalism.

The initial state of an object o of class C with actual class parameter values \bar{v} is denoted $init_{o:C(\bar{v})}$ and is defined by

$$init_{o:C(\bar{v})} \triangleq \mathbf{object}(Id : o, Cl : C, Pr : init_C, Lvar : \emptyset, Flds : a, Cnt : 0, Mtd : init)$$

where a is the initial state of the object fields given by [**this** $\mapsto o, Par_C \mapsto \bar{v}, Flds_C \mapsto \bar{d}$]. Here Par_C , $Flds_C$, and $init_C$, represent the class parameters, the fields, and the initialization code of C , respectively. The class parameters Par_C are initialized by the actual parameters \bar{v} , the fields $Flds_C$ are initialized by default values \bar{d} (of the appropriate types), and the initial code is ready to be executed with an empty local state.

A system is given by a set of self-contained classes \overline{Cl} , including a class $Main$, without class parameters, used to generate the initial object $init_{main:Main(\varepsilon)}$. The initial configuration of a system is defined by

$$init_{\overline{Cl}} \triangleq \overline{Cl} \quad init_{main:Main(\varepsilon)} \quad \mathbf{hist}(\varepsilon)$$

The operational rules are summarized in Fig. 4. The rules for skip, assignment, initialized variable declarations, if- and while-statements are standard. Note that $(a; l)$ represents the total object state, composed by a , the state of the fields/class parameters, and l , the state of the local variables/parameters of the method. Lookup of a variable if left to right, i.e., l is tried before a . Expressions e without side-effects are evaluated by a semantic function depending on the total state, i.e., $eval(e, (a; l))$.

Method invocation is captured by the rule **call**. The generated future identity $ft(o, n)$ is globally unique (assuming the *next* function is producing locally unique values). The future unit itself is not generated yet; it will be generated by return from the called method.

If there is no active process in an object, denoted $Pr : empty$, a method call is selected for execution by rule **method**. The invocation message is consumed by this rule, and the future identity of the call is assigned to the implicit parameter **destiny**. Method execution is completed by rule **return**, and a future value is fetched by rule **query**. A query can only succeed if the appropriate future unit is generated. A future unit appears in the configuration when resolved by rule **return**, which means that a query statement blocks until the future is resolved. Remark that rule **query** does not remove the future unit from the configuration, which allows several processes to fetch the value of the same future.

In rule **new**, the new object gets a unique identity, $ob(C, n)$, given by that of the generating object and a counter, the actual class parameters are evaluated, and the initialization is performed. The given language fragment may be extended with constructs for inter object process control and suspension, e.g., by using the *ABS* approach of [18].

skip : $\mathbf{object}(Id : o, Pr : (\mathbf{skip}; s)) \longrightarrow \mathbf{object}(Id : o, Pr : s)$
assign : $\mathbf{object}(Id : o, Pr : (v := e; s), Lvar : l, FlDs : a)$
 \longrightarrow
 $\mathbf{if } v \text{ in } l \text{ then } \mathbf{object}(Id : o, Pr : s, Lvar : l[v \mapsto \mathit{eval}(e, (a; l))], FlDs : a)$
 $\mathbf{else } \mathbf{object}(Id : o, Pr : s, Lvar : l, FlDs : a[v \mapsto \mathit{eval}(e, (a; l))])$
init : $\mathbf{object}(Id : o, Pr : (\overline{Tv} := e; s), Lvar : l, FlDs : a)$
 \longrightarrow
 $\mathbf{object}(Id : o, Pr : (\overline{Tv}; \overline{v} := \overline{e}; s), Lvar : l, FlDs : a)$
if-else : $\mathbf{object}(Id : o, Pr : (\mathit{if } e \text{ then } s_1 \text{ else } s_2 \text{ fi}; s), Lvar : l, FlDs : a)$
 \longrightarrow
 $\mathbf{if } \mathit{eval}(e, (a; l)) \text{ then } \mathbf{object}(Id : o, Pr : (s_1; s), Lvar : l, FlDs : a)$
 $\mathbf{else } \mathbf{object}(Id : o, Pr : (s_2; s), Lvar : l, FlDs : a)$
while : $\mathbf{object}(Id : o, Pr : (\mathit{while } e \text{ do } s_1 \text{ od}; s), Lvar : l, FlDs : a)$
 \longrightarrow
 $\mathbf{object}(Id : o, Pr : (\mathit{if } e \text{ then } s_1; \mathit{while } e \text{ do } s_1 \text{ od fi}; s), Lvar : l, FlDs : a)$
new : $\mathbf{hist}(h) \mathbf{class}(Id : C, Cnt : n)$
 $\mathbf{object}(Id : o, Pr : (v := \mathit{new } C(\overline{e}); s), Lvar : l, FlDs : a)$
 \longrightarrow
 $\mathbf{hist}(h \cdot \langle o \xrightarrow{\mathit{new}} \mathit{ob}(C, n), C, \mathit{eval}(\overline{e}, (a; l)) \rangle) \mathbf{class}(Id : C, Cnt : \mathit{next}(n))$
 $\mathbf{object}(Id : o, Pr : (v := \mathit{ob}(C, n); s), Lvar : l, FlDs : a)$
 $\mathit{init}_{\mathit{ob}(C, n):C}(\mathit{eval}(\overline{e}, (a; l)))$
call : $\mathbf{hist}(h) \mathbf{object}(Id : o, Pr : (fr := v!m(\overline{e}); s), Lvar : l, FlDs : a, Cnt : n)$
 \longrightarrow
 $\mathit{MSG} \mathbf{hist}(h \cdot \mathit{MSG})$
 $\mathbf{object}(Id : o, Pr : (fr := \mathit{ft}(o, n); s), Lvar : l, FlDs : a, Cnt : \mathit{next}(n))$
method : $\langle o' \rightarrow o, u, m, \overline{v} \rangle \mathbf{hist}(h) \mathbf{class}(Id : c, Mtds : (q (m, \overline{p}, l, s)))$
 $\mathbf{object}(Id : o, Cl : c, Pr : \mathit{empty}, FlDs : a)$
 \longrightarrow
 $\mathbf{hist}(h \cdot \langle o' \rightarrow o, u, m, \overline{v} \rangle) \mathbf{class}(Id : c, Mtds : (q (m, \overline{p}, l, s)))$
 $\mathbf{object}(Id : o, Cl : c, Pr : s, Lvar : l[\overline{p} \mapsto \overline{v}][\mathit{destiny} \mapsto u], FlDs : a, Mtd : m)$
return : $\mathbf{hist}(h) \mathbf{object}(Id : o, Pr : \mathbf{return } e, Lvar : l, FlDs : a, Mtd : m)$
 \longrightarrow
 $\mathbf{hist}(h \cdot \langle \leftarrow o, \mathit{eval}(\mathit{destiny}, l), m, \mathit{eval}(e, (a; l)) \rangle)$
 $\mathbf{fut}(Id : \mathit{eval}(\mathit{destiny}, l), Val : \mathit{eval}(e, (a; l)))$
 $\mathbf{object}(Id : o, Pr : \mathit{empty}, FlDs : a)$
query : $\mathbf{hist}(h) \mathbf{fut}(Id : u, Val : d) \mathbf{object}(Id : o, Pr : (v := e?; s), Lvar : l, FlDs : a)$
 \longrightarrow
 $\mathbf{hist}(h \cdot \langle o \leftarrow u, d \rangle) \mathbf{fut}(Id : u, Val : d)$
 $\mathbf{object}(Id : o, Pr : (v := d; s), Lvar : l, FlDs : a)$
 $\mathbf{if } \mathit{eval}(e, (a; l)) = u$

Fig. 4. Operational rules, using the standard rewriting logic convention that irrelevant attributes may be omitted in a rule. Variables are denoted by single characters (the uniform naming convention is left implicit), $(a; l)$ represents the total object state, and $a[v \mapsto d]$ is the state a updated by binding the variable v to the data value d . The eval function evaluates an expression in a given state, and in is used for testing domain membership. In rule **call**, MSG denotes $\langle o \rightarrow \mathit{eval}(v, (a; l)), \mathit{ft}(o, n), m, \mathit{eval}(\overline{e}, (a; l)) \rangle$.

4.2 Semantic Properties

Semantic properties are stated by means of notions of validity. We define *global validity* (denoted \models) and *local validity* with respect to a class C (denoted \models_C). A global object system initiated by a configuration $\mathit{init}_{\overline{Cl}}$ is said to satisfy a global

invariant property $I(h)$, if the global history h of any reachable configuration G satisfies $I(h)$:

$$\overline{Cl} \models I(h) \triangleq \forall G. \text{init}_{\overline{Cl}} \longrightarrow^* G \wedge G.\text{hist} = h \Rightarrow I(h)$$

where \longrightarrow^* denotes the transitive and reflexive extension of the transition relation, lifted to configurations, and where $G.\text{hist}$ extracts the history of the configuration G .

Similarly, an object system initiated by a configuration $\text{init}_{\overline{Cl}}$ is said to satisfy a C -local invariant property $I(h)$ if every object o of class C in any reachable configuration G satisfies $I(h/o)$, i.e., the projection from global history to the object o :

$$\overline{Cl} \models_C I(h) \triangleq \forall G, o. \text{init}_{\overline{Cl}} \longrightarrow^* G \wedge G.\text{hist} = h \wedge o \in G.\text{obj} \wedge G[o].\text{class} = C \Rightarrow I(h/o)$$

where $G.\text{obj}$ extracts the object identities from the objects in the configuration G .

We next provide notions of global and local well-formedness for global histories. We first introduce some notation and functions used in defining wellformed histories. For sequences a and b , let $a \mathbf{ew} x$ denote that x is the last element of a , $\text{agree}(a)$ denote that all elements (if any) are equal, and $a \leq b$ denote that a is a prefix of b . Let $[x_1, x_2, \dots, x_i]$ denote the sequence of x_1, x_2, \dots, x_i for $i > 0$ (allowing repeated parts [...]*). Functions for event decomposition are lifted to sequences in the standard way, ignoring events for which the decomposition is not defined, e.g., $_.\text{result} : \text{Seq}[\text{Ev}] \rightarrow \text{Seq}[\text{Data}]$.

Functions may extract information from the history. In particular, we define $\text{oid} : \text{Seq}[\text{Ev}] \rightarrow \text{Set}[\text{Obj}]$ extracting all object identities occurring in a history, as follows:

$$\begin{aligned} \text{oid}(\varepsilon) &\triangleq \{\text{main}\} & \text{oid}(h \cdot \gamma) &\triangleq \text{oid}(h) \cup \text{oid}(\gamma) \\ \text{oid}(\langle o \rightarrow o', u, m, \bar{e} \rangle) &\triangleq \{o, o'\} \cup \text{oid}(\bar{e}) & \text{oid}(\langle o' \rightarrow o, u, m, \bar{e} \rangle) &\triangleq \{o, o'\} \cup \text{oid}(\bar{e}) \\ \text{oid}(\langle \leftarrow o, u, m, e \rangle) &\triangleq \{o\} \cup \text{oid}(e) & \text{oid}(\langle o \leftarrow, u, e \rangle) &\triangleq \{o\} \cup \text{oid}(e) \\ \text{oid}(\langle o \xrightarrow{\text{new}} o', C, \bar{e} \rangle) &\triangleq \{o, o'\} \cup \text{oid}(\bar{e}) & & \end{aligned}$$

where $\gamma : \text{Ev}$, and $\text{oid}(\bar{e})$ returns the set of object identifiers occurring in the expression list \bar{e} . The function $\text{fid} : \text{Seq}[\text{Ev}] \rightarrow \text{Set}[\text{Fid}]$ extracts future identities from a history:

$$\begin{aligned} \text{fid}(\varepsilon) &\triangleq \emptyset & \text{fid}(h \cdot \gamma) &\triangleq \text{fid}(h) \cup \text{fid}(\gamma) \\ \text{fid}(\langle o \rightarrow o', u, m, \bar{e} \rangle) &\triangleq \{u\} & \text{fid}(\langle o' \rightarrow o, u, m, \bar{e} \rangle) &\triangleq \{u\} \cup \text{fid}(\bar{e}) \\ \text{fid}(\langle \leftarrow o, u, m, e \rangle) &\triangleq \emptyset & \text{fid}(\langle o \leftarrow, u, e \rangle) &\triangleq \text{fid}(e) \\ \text{fid}(\langle o \xrightarrow{\text{new}} o', C, \bar{e} \rangle) &\triangleq \text{fid}(\bar{e}) & & \end{aligned}$$

where $\gamma : \text{Ev}$, and $\text{fid}(\bar{e})$ returns the set of future identities occurring in the expression list \bar{e} . For a global history h , the function $\text{fid}(h)$ returns all future identities on h , and for a local history h/o , the function $\text{fid}(h/o)$ returns the futures generated by o or received as parameters. At last, h/u abbreviates the projection of history h to the set $\{\gamma \mid \gamma.\text{future} = u\}$, i.e., all events with future u .

Definition 4. (Wellformed histories) Let $h : \text{Seq}[E\mathcal{V}]$ be a history of a global object system S . The well-formedness predicate $wf : \text{Seq}[E\mathcal{V}] \rightarrow \text{Bool}$ is defined by:

$$\begin{aligned}
wf(\varepsilon) &\triangleq \text{true} \\
wf(h \cdot \langle o \rightarrow o', u, m, \bar{e} \rangle) &\triangleq wf(h) \wedge o \neq \text{null} \wedge u \notin \text{fid}(h) \cup \text{fid}(\bar{e}) \\
wf(h \cdot \langle o' \rightarrow o, u, m, \bar{e} \rangle) &\triangleq wf(h) \wedge o \neq \text{null} \wedge h/u = [\langle o' \rightarrow o, u, m, \bar{e} \rangle] \\
wf(h \cdot \langle \leftarrow o, u, m, e \rangle) &\triangleq wf(h) \wedge h/u \text{ **ew** } \langle _ \rightarrow o, u, m, _ \rangle \\
wf(h \cdot \langle o \leftarrow, u, e \rangle) &\triangleq wf(h) \wedge u \in \text{fid}(h/o) \wedge \text{agree}((h/u).\text{result}) \cdot e \\
wf(h \cdot \langle o \xrightarrow{\text{new}} o', C, \bar{e} \rangle) &\triangleq wf(h) \wedge o \neq \text{null} \wedge o' \neq \text{null} \wedge o' \notin \text{oid}(h) \cup \text{oid}(\bar{e})
\end{aligned}$$

It follows directly that a wellformed global history satisfies the communication order pictured in Fig. 3, i.e.,

$$\forall u. \exists o, o', m, \bar{e}, e. \\
h/u \leq [\langle o' \rightarrow o, u, m, \bar{e} \rangle, \langle o' \rightarrow o, u, m, \bar{e} \rangle, \langle \leftarrow o, u, m, e \rangle, [\langle _ \leftarrow, u, e \rangle]^*]$$

Also, it ensures the uniqueness of object identifiers and future identities. We can prove that the operational semantics guarantees well-formedness:

Lemma 1. *The global history h of a global object system S obtained by the given operational semantics, is wellformed, i.e., $\models wf(h)$ where $wf(h)$ is strengthened by the two conditions $\text{fid}(h) \subseteq (h/\rightarrow).\text{future}$ and $\text{oid}(h) - \text{null} \subseteq (h/\xrightarrow{\text{new}}).\text{callee}$.*

The two conditions ensure that a history may not refer to object and future identities before generated by creation and invocation events, respectively. This lemma follows by induction over the number of rule applications.

Well-formedness of a local history for an object o , denoted $wf_o(h)$, is defined as in Def. 4, except that the last conjunct of the case $\langle o' \rightarrow o, u, m, \bar{e} \rangle$ only holds for self calls, i.e., where o and o' are equal. For local well-formedness, the conjunct is therefore weakened to $o = o' \Rightarrow h/u = [\langle o' \rightarrow o, u, m, \bar{e} \rangle]$. If h is a wellformed global history, it follows immediately that each projection h/o is locally wellformed, i.e.,

$$wf(h) \Rightarrow wf_o(h/o)$$

5 Program Verification

The communication history abstractly captures the system state at any point in time [11, 12]. Partial correctness properties of a system may thereby be specified by finite initial segments of its communication histories. A *history invariant* is a predicate over the communication history, which holds for all finite sequences in the (prefix-closed) set of possible histories, expressing safety properties [5]. In this section we present a framework for compositional reasoning about object systems, establishing an invariant over the global history from invariants over

the local histories of each object. Since the local object histories are disjoint with our four event semantics, it is possible to reason locally about each object. In particular, the history updates of the operational semantics affect the local history of the active object only, and can be treated simply as an assignment to the local history. The local history is not effected by the environment, and interference-free reasoning is then possible. Correspondingly, the reasoning framework consists of two parts: A proof system for local (class-based) reasoning, and a rule for composition of object specifications.

5.1 Local Reasoning

Pre- and postconditions to method definitions are in our setting used to establish a *class invariant*. The class invariant must hold after initialization of all class instances and must be maintained by all methods, serving as a contract for the different methods: A method implements its part of the contract by ensuring that the invariant holds upon termination, assuming that it holds when the method starts execution. A class invariant establishes a *relationship between the internal state and the observable behavior of class instances*. The internal state reflects the values of the fields, and the observable behavior is expressed as potential communication histories. A *user-provided invariant* $I(\bar{w}, \mathcal{H})$ for a class C is a predicate over the fields \bar{w} , the read-only parameters $\bar{c}p$ and **this**, in addition to the local history \mathcal{H} which is a sequence of events generated by **this**. The proof system for class-based verification is formulated within *dynamic logic* as used by the KeY framework [7], facilitating class invariant verification by considering each method independently. The dynamic logic formulation suggests that the proof system is suitable for an implementation in the KeY framework.

Dynamic logic provides a structured way to describe program behavior by an integration of programs and assertions within a single language. The formula $[s]\phi$ expresses the precondition of s with ϕ as postcondition. The formula $\psi \Rightarrow [s]\phi$ express partial correctness properties: if statement s is executed in a state where ψ holds and the execution terminates, then ϕ holds in the final state. The formula is verified by a symbolic execution of s , where state modifications are handled by the *update* mechanism [7]. A dynamic formula $[s_1; s]\phi$ is equal to $[s_1][s]\phi$. A dynamic formula $[v := e; s]\phi$, i.e., where an assignment is the first statement, reduces to $\{v := e\}[s]\phi$, where $\{v := e\}$ is an update. We assume that expressions e can be evaluated within the assertion language. Updates can only be *applied* on formulas without programs, which means that updates on a formula $[s]\phi$ are accumulated and *delayed* until the symbolic execution of s is complete. Update application $\{v := t\}e$, on an expression e , evaluates to the substitution e_t^v , replacing all free occurrences of v in e by t . The parallel update $\{v_1 := e_1 || \dots || v_n := e_n\}$, for disjoint variables v_1, \dots, v_n , represents an accumulated update, and the application of a parallel update leads to a simultaneous substitution. For an update U , we have $U(\phi_1 \wedge \phi_2) = U\phi_1 \wedge U\phi_2$. A *sequent* $\psi_1, \dots, \psi_n \vdash \phi_1, \dots, \phi_m$ contains assumptions ψ_1, \dots, ψ_n , and formulas ϕ_1, \dots, ϕ_m to be proved. The sequent is *valid* if at least one formula ϕ_i follows from the assumptions, and it can be interpreted as $\psi_1 \wedge \dots \wedge \psi_n \Rightarrow \phi_1 \vee \dots \vee \phi_m$.

$$\begin{array}{c}
\text{invoc} \frac{\vdash \forall u. \{ \mathcal{H} := \mathcal{H} \cdot \langle \text{this} \rightarrow v, u, m, \bar{e} \rangle \mid fr := u \} [s] \phi}{\vdash [fr := v!m(\bar{e}); s] \phi} \\
\text{fetch} \frac{\vdash \forall v'. \{ \mathcal{H} := \mathcal{H} \cdot \langle \text{this} \leftarrow, e, v' \rangle \mid v := v' \} ([s] \phi \wedge \exists \bar{w}. I(\bar{w}, \mathcal{H}))}{\vdash [v := e?; s] \phi} \\
\text{new} \frac{\vdash \forall v'. \{ \mathcal{H} := \mathcal{H} \cdot \langle \text{this} \xrightarrow{\text{new}} v', C, \bar{e} \rangle \mid v := v' \} [s] \phi}{\vdash [v := \mathbf{new} C(\bar{e}); s] \phi}
\end{array}$$

Fig. 5. Dynamic logic rules for method invocation, future query and object creation. $I(\bar{w}, \mathcal{H})$ is the class invariant.

In order to verify a class invariant $I(\bar{w}, \mathcal{H})$, we must prove that the invariant is established by the initialization code and maintained by all method definitions in C , assuming well-formedness of the local history. For a method definition $m(\bar{x})\{s; \mathbf{return} e\}$ in C , this amounts to a proof of the sequent:

$$\vdash (wf_{\text{this}}(\mathcal{H}) \wedge I(\bar{w}, \mathcal{H}) \Rightarrow [\mathcal{H} := \mathcal{H} \cdot \langle \text{caller} \rightarrow \text{this}, \text{destiny}, m, \bar{x} \rangle; s; \mathcal{H} := \mathcal{H} \cdot \langle \leftarrow \text{this}, \text{destiny}, m, e \rangle](wf_{\text{this}}(\mathcal{H}) \Rightarrow I(\bar{w}, \mathcal{H}))$$

Here, the method body is extended with a statement for extending the history with the invocation reaction event, and the **return** statement is treated as a history extension. Dynamic logic rules for method invocation, future query, and object creation, can be found in Fig. 5. When invoking a method, the update in the premise of rule `invoc` captures the history extension and the generation of a fresh future identity u . Similarly, the update in rule `fetch` captures the history extension and the assignment of a fresh value to v , where the well-formedness assumptions ensure that all values received from the same future are equal. The update in the premise of rule `new` captures the history extension and the generation of a fresh object identity v' , and the universal quantifier reflects non-determinism. The prime is needed here since v may occur in \bar{e} . The query rule insists that the class invariant holds for local history, ignoring the field values of the current state, as discussed in the soundness proof. Assignments are analyzed as explained above, and rules for **skip** and conditionals are standard. We refer to Din et al. for further details [19].

The rules for the rest of the *ABS* statements can be defined as substitution rules introduced in Fig. 6. For instance, `[skip; s]φ` can be rewritten to `[s]φ`. In rule `declNit` and `declNoNit` v' is needed since the postcondition may talk about a field with the same name v . If-statements without an else-branch are as usual.

5.2 Soundness

The reasoning system for statements in dynamic logic is sound if any provable property is valid, i.e.,

$$\vdash \psi \Rightarrow [s] \phi \Rightarrow \models \psi \Rightarrow [s] \phi$$

skip	$[\mathbf{skip}; s]\phi = [s]\phi$
assign	$[v := e; s]\phi = \{v := e\} [s]\phi$
declNnit	$[T \ v = e; s]\phi = [v' := e; s_{v'}^v]\phi$
declNolnit	$[T \ v; s]\phi = [v' := \mathbf{default}_T; s_{v'}^v]\phi$
ifElse	$[\mathbf{if} \ b \ \mathbf{then} \ s_1 \ \mathbf{else} \ s_2 \ \mathbf{fi}; s]\phi = \mathbf{if} \ b \ \mathbf{then} \ [s_1; s]\phi \ \mathbf{else} \ [s_2; s]\phi$
while	$[\mathbf{while} \ b \ \mathbf{do} \ s' \ \mathbf{od}; s]\phi = \mathbf{if} \ b \ \mathbf{then} \ (\exists w. I(\bar{w}, \mathcal{H})) \wedge$ $([s'; \mathbf{while} \ b \ \mathbf{do} \ s' \ \mathbf{od}; s]\phi) \ \mathbf{else} \ [s]\phi$

Fig. 6. Semantical definitions for standard *ABS* statements. Here ϕ is the postcondition, s is the remaining program yet to be executed, primes denote fresh variables, $s_{v'}^v$ is s with all (free) occurrences of v replaced by v' , and $\mathbf{default}_T$ is the default value defined for type T .

Validity of a dynamic logic formula, denoted $\models \psi \Rightarrow [s]\phi$, is defined by means of the operational semantics. We base the semantics on the operational semantics above, as given by unlabeled transitions of the form $G_1 \rightarrow G_2$.

Note that each rule is local to one object, and we write $G_1 \xrightarrow{o:s} G_2$ to indicate an execution involving only object o such that exactly the statement (list) s has been executed by o . And we write $G_1 \rightrightarrows G_2$ if o executes s while other objects may execute.

Definition 5 (Explicit execution step).

$$G_1 \xrightarrow{o:s} G_2 \triangleq G_1 \longrightarrow^* G_2 \wedge G_1[o].Pr = s; G_2[o].Pr$$

$$G_1 \rightrightarrows G_2 \triangleq G_1 \xrightarrow{o:s} G_2 \wedge \forall o'. o' \neq o \Rightarrow G_1[o'] = G_2[o']$$

expressing one or more transitions from the configuration G_1 to G_2 such that o executes s , with or without, respectively, interleaved execution by other objects. The notation $G[o]$ denotes the object o of the configuration G .

We consider pre- and postconditions over local states and the local history. Such an assertion can be evaluated in a state defining values for attributes (of the appropriate class), parameters and local variables (of the method) and the local history. We let $\llbracket \psi \xrightarrow{o:s} \phi \rrbracket_{G,o}$ express that if the condition ψ holds for object o before execution of s by the object in configuration G , then ϕ holds for o after the execution. As above, we let $\xrightarrow{o:s}$ express local execution by o , and \rightrightarrows execution by o interleaved with other objects:

Definition 6 (Validity of pre/post-conditions over execution steps).

$$\llbracket \psi \xrightarrow{o:s} \phi \rrbracket_{G,o} \triangleq \forall G', \bar{z}. wf(G'.hist) \wedge G \xrightarrow{o:s} G' \wedge loc(G,o)[\psi] \Rightarrow loc(G',o)[\phi]$$

$$\llbracket \psi \rightrightarrows \phi \rrbracket_{G,o} \triangleq \forall G', \bar{z}. wf(G'.hist) \wedge G \rightrightarrows G' \wedge loc(G,o)[\psi] \Rightarrow loc(G',o)[\phi]$$

where \bar{z} is the list of auxiliary variables in ψ and/or ϕ , not bound by G nor G' . Here $loc(G, o)$ denotes the local state of object o , as derived from the global state G . The function $loc : \text{Config} \times \text{Oid} \rightarrow \text{State}$ is defined by

$$loc(G, o) \triangleq (G[o].Flds; G[o].Lvar) + [\mathcal{H} \mapsto (G.hist)/o]$$

where the resulting \mathcal{H} ranges over local histories (i.e., in the alphabet of o), and where this is bound to o in G as explained earlier. Thus the extraction is made by taking the state of object o and adding the history localized to o . We let $loc(G, o)[\psi]$ denote the value of ψ in state $loc(G, o)$.

It follows that local reasoning suffices for local pre/post-conditions, in the sense that when reasoning about one object in our system, one may ignore the activity of other objects.

Lemma 2. $\llbracket \psi \xrightarrow{o:s} \phi \rrbracket_{G,o}$ is the same as $\llbracket \psi \rightrightarrows \phi \rrbracket_{G,o}$

The lemma follows by induction on the length of executions, and the fact that $loc(G, o)$ for any G is not affected by execution steps by other objects than o , since remote access to fields is not allowed in our language and since h/o only contains events generated by o .

In our setting, we may understand a sequent by means of the $\xrightarrow{o:s}$ relation, letting a dynamic logic subformula depend on a given pre-configuration G and object o .

Definition 7 (Validity of dynamic logic sequents).

$$\begin{aligned} \models \psi_1, \dots, \psi_n \vdash \phi_1, \dots, \phi_m &\triangleq \forall G, o. wf(G.hist) \Rightarrow \llbracket \psi_1 \wedge \dots \wedge \psi_n \Rightarrow \phi_1 \vee \dots \vee \phi_m \rrbracket_{G,o} \\ \llbracket [s]\phi \rrbracket_{G,o} &\triangleq \llbracket true \xrightarrow{o:s} \phi \rrbracket_{G,o} \\ \llbracket e \rrbracket_{G,o} &\triangleq loc(G, o)[e] \\ \llbracket \{v := t\}e \rrbracket_{G,o} &\triangleq \llbracket e \xrightarrow{v} \rrbracket_{G,o} \\ \llbracket \psi \wedge \phi \rrbracket_{G,o} &\triangleq \llbracket \psi \rrbracket_{G,o} \wedge \llbracket \phi \rrbracket_{G,o} \\ \llbracket \psi \vee \phi \rrbracket_{G,o} &\triangleq \llbracket \psi \rrbracket_{G,o} \vee \llbracket \phi \rrbracket_{G,o} \\ \llbracket \psi \Rightarrow \phi \rrbracket_{G,o} &\triangleq \llbracket \psi \rrbracket_{G,o} \Rightarrow \llbracket \phi \rrbracket_{G,o} \\ \llbracket U(\psi \wedge \phi) \rrbracket_{G,o} &\triangleq \llbracket U\psi \rrbracket_{G,o} \wedge \llbracket U\phi \rrbracket_{G,o} \\ \llbracket U(\psi \vee \phi) \rrbracket_{G,o} &\triangleq \llbracket U\psi \rrbracket_{G,o} \vee \llbracket U\phi \rrbracket_{G,o} \\ \llbracket U(\psi \Rightarrow \phi) \rrbracket_{G,o} &\triangleq \llbracket U\psi \rrbracket_{G,o} \Rightarrow \llbracket U\phi \rrbracket_{G,o} \end{aligned}$$

where e is a formula without the dynamic logic operators, and the equations for updates are as given earlier. The application of a parallel update U , for instance, $\{v_1 := t_1 \mid \dots \mid v_n := t_n\}$, is for short written as $\{v := t\}$. It follows from the definition that

$$\llbracket \psi \Rightarrow [s]\phi \rrbracket_{G,o} = \llbracket \psi \xrightarrow{o:s} \phi \rrbracket_{G,o}$$

Here o is the executing object and the object on which ψ and ϕ are interpreted. Thus the formula is valid if for any object o executing s , the postcondition holds in the poststate, provided the precondition holds in the prestate. In dynamic

logic the prestate given by G and o is fixed for the whole sequent, and therefore the meaning of the individual operators is given in the context of G and o .

We verify an invariant $I(\bar{w}, \mathcal{H})$ for a class C by showing that $I(\bar{w}, \mathcal{H})$ is established by the initialization of C , i.e. $init_C$, and is maintained by all methods in C , assuming local well-formedness. The rule is:

$$\begin{array}{l}
\vdash \exists \bar{w}. I(\bar{w}, \mathcal{H} \cdot \gamma) \Rightarrow \exists \bar{w}. I(\bar{w}, \mathcal{H}) \\
\vdash \mathcal{H} = \varepsilon \Rightarrow [init_C](wf_{this}(\mathcal{H}) \Rightarrow I(\bar{w}, \mathcal{H})) \\
\vdash (wf_{this}(\mathcal{H}) \wedge I(\bar{w}, \mathcal{H})) \Rightarrow [body_{C,m}](wf_{this}(\mathcal{H}) \Rightarrow I(\bar{w}, \mathcal{H})), \text{ for all methods } m \text{ in } C \\
\text{class} \frac{}{\vdash_C \exists \bar{w}. I(\bar{w}, \mathcal{H})}
\end{array}$$

where $body_{C,m}$ denotes the body s of method m of C augmented with effects on the local history reflecting the start and end of the method, namely

$$\mathcal{H} := \mathcal{H} \cdot \langle \text{caller} \rightarrow \text{this}, \text{destiny}, m, \bar{x} \rangle; s; \mathcal{H} := \mathcal{H} \cdot \langle \leftarrow \text{this}, \text{destiny}, m, e \rangle$$

Lemma 3. *Reasoning about statements is sound:*

$$\vdash \psi \Rightarrow [s]\phi \Rightarrow \models \psi \Rightarrow [s]\phi$$

Theorem 1. *The proof system for reasoning about classes is sound:*

$$\vdash_C I(\mathcal{H}) \Rightarrow \models_C I(\mathcal{H})$$

5.2.0.1 Proof of lemma 3. We focus on the rules for statements involving futures and object generation, and consider therefore the rules `invoc`, `fetch` and `new`, as given in Fig. 5. The axioms given in Fig. 6 represent standard statements not involving futures, and we omit the soundness proof of these.

5.2.0.2 Asynchronous method call statement

We prove that the `invoc` rule preserves validity. The validity of the conclusion is $\models [fr := v!m(\bar{e}); s]\phi$. Consider now a given G and o , and let ϕ' denote $[s]\phi$. According to Def. 7, the validity can be written as

$$wf(G.hist) \Rightarrow \llbracket true \xrightarrow{o:fr:=v!m(\bar{e})} \phi' \rrbracket_{G,o}$$

which by Def. 6 is

$$\forall G', \bar{z}. wf(G.hist) \wedge wf(G'.hist) \wedge G \xrightarrow{o:fr:=v!m(\bar{e})} G' \Rightarrow loc(G', o)[\phi']$$

By the operational semantics of `call` and `assign`, we have that G' is G with `MSG` and $G'.hist = G.hist \cdot \text{MSG}$, where `MSG` denotes $\langle o \rightarrow loc(G, o)[v], ft(o, n), m, loc(G, o)[\bar{e}] \rangle$, and such that the object state $G'[o].l$ is $(G[o].l)[fr \mapsto ft(o, n)]$ if $fr \in G[o].l$, and otherwise $G'[o].a$ is $(G[o].a)[fr \mapsto ft(o, n)]$. Here n is the counter value of $G[o]$ (same as in $G'[o]$). Other parts of the object state are unchanged.

Thus $loc(G', o)[\phi']$ can be reduced to $loc(G, o)[\phi' \xrightarrow{fr, \mathcal{H}}_{ft(o, n), \mathcal{H} \cdot \langle \text{this} \rightarrow v, ft(o, n), m, \bar{e} \rangle}]$ since $loc(G, o)[\langle \text{this} \rightarrow v, ft(o, n), m, \bar{e} \rangle] = \text{MSG}$, and it suffices to prove

$$\forall \bar{z}. wf(G.hist) \Rightarrow loc(G, o)[\phi' \xrightarrow{fr, \mathcal{H}}_{ft(o, n), \mathcal{H} \cdot \langle \text{this} \rightarrow v, ft(o, n), m, \bar{e} \rangle}]$$

The validity of the premise is

$$\models \forall u. \{ \mathcal{H} := \mathcal{H} \cdot \langle \text{this} \rightarrow v, u, m, \bar{e} \rangle \mid fr := u \} \phi'$$

which by Def. 7 is

$$\forall \bar{z}, u. wf(G.hist) \Rightarrow loc(G, o)[\phi' \xrightarrow{fr, \mathcal{H}}_{u, \mathcal{H} \cdot \langle \text{this} \rightarrow v, u, m, \bar{e} \rangle}]$$

Clearly this is sufficient to ensure validity of the conclusion, since the universal quantifier on u covers the value given by $ft(o, n)$.

5.2.0.3 Query statement

We prove that the fetch rule preserves validity. The validity of the conclusion is $\models [v := e?; s]\phi$. Consider now a given G and o , and let ϕ' denote $[s]\phi$. According to Def. 7, the validity can be written as

$$wf(G.hist) \Rightarrow \llbracket true \xrightarrow{o:v:=e?} \phi' \rrbracket_{G, o}$$

which by Def. 6 is

$$\forall G', \bar{z}. wf(G.hist) \wedge wf(G'.hist) \wedge G \xrightarrow{o:v:=e?} G' \Rightarrow loc(G', o)[\phi']$$

By the operational semantics of `query` and `assign`, we have that G' is G with `MSG` and $G'.hist = G.hist \cdot \text{MSG}$, where `MSG` denotes $\langle o \leftarrow, loc(G, o)[e], d \rangle$ and such that the object state $G'[o].l$ is $(G[o].l)[v \mapsto d]$ if $v \in G[o].l$, and otherwise $G'[o].a$ is $(G[o].a)[v \mapsto d]$. Other parts of the object state are unchanged.

Thus $loc(G', o)[\phi']$ can be reduced to $loc(G, o)[\phi' \xrightarrow{v, \mathcal{H}}_{d, \mathcal{H} \cdot \langle \text{this} \leftarrow, e, d \rangle}]$ since $loc(G, o)[\langle \text{this} \leftarrow, e, d \rangle] = \text{MSG}$, and it suffices to prove

$$\forall \bar{z}. wf(G.hist) \Rightarrow loc(G, o)[\phi' \xrightarrow{v, \mathcal{H}}_{d, \mathcal{H} \cdot \langle \text{this} \leftarrow, e, d \rangle}]$$

The validity of the premise is

$$\models \forall v'. \{ \mathcal{H} := \mathcal{H} \cdot \langle \text{this} \leftarrow, e, v' \rangle \mid v := v' \} (\phi' \wedge \exists \bar{w}. I(\bar{w}, \mathcal{H}))$$

which by Def. 7 is

$$\forall \bar{z}, v'. wf(G.hist) \Rightarrow loc(G, o)[(\phi' \wedge \exists \bar{w}. I(\bar{w}, \mathcal{H})) \xrightarrow{v, \mathcal{H}}_{v', \mathcal{H} \cdot \langle \text{this} \leftarrow, e, v' \rangle}]$$

Clearly this is sufficient to ensure validity of the conclusion, since the universal quantifier on v' covers the value given by d . Note that the invariant is not required here.

5.2.0.4 Object creation statement

We prove that the new rule preserves validity. The validity of the conclusion is $\models [v := \text{new } C(\bar{e}); s]\phi$. Consider now a given G and o , and let ϕ' denote $[s]\phi$. According to Def. 7, the validity can be written as

$$wf(G.hist) \Rightarrow \llbracket true \xrightarrow{o:v:=\text{new } C(\bar{e})} \phi' \rrbracket_{G,o}$$

which by Def. 6 is

$$\forall G', \bar{z}. wf(G.hist) \wedge wf(G'.hist) \wedge G \xrightarrow{o:v:=\text{new } C(\bar{e})} G' \Rightarrow loc(G', o)[\phi']$$

By the operational semantics of `new` and `assign`, we have that G' is G with `MSG` and $G'.hist = G.hist \cdot \text{MSG}$, where `MSG` denotes $\langle o \xrightarrow{\text{new}} ob(C, n), C, loc(G, o)[\bar{e}] \rangle$ and such that the object state $G'[o].l$ is $(G[o].l)[v \mapsto ob(C, n)]$ if $v \in G[o].l$, and otherwise $G'[o].a$ is $(G[o].a)[v \mapsto ob(C, n)]$. Here n is the counter value of $G[C]$ (same as in $G'[C]$). Other parts of the object state are unchanged.

Thus $loc(G', o)[\phi']$ can be reduced to $loc(G, o)[\phi' \xrightarrow{v, \mathcal{H}}_{ob(C, n), \mathcal{H} \cdot \langle \text{this} \xrightarrow{\text{new}} ob(C, n), C, \bar{e} \rangle}]$ since $loc(G, o)[\langle \text{this} \xrightarrow{\text{new}} ob(C, n), C, \bar{e} \rangle] = \text{MSG}$, and it suffices to prove

$$\forall \bar{z}. wf(G.hist) \Rightarrow loc(G, o)[\phi' \xrightarrow{v, \mathcal{H}}_{ob(C, n), \mathcal{H} \cdot \langle \text{this} \xrightarrow{\text{new}} ob(C, n), C, \bar{e} \rangle}]$$

The validity of the premise is

$$\models \forall v'. \{ \mathcal{H} := \mathcal{H} \cdot \langle \text{this} \xrightarrow{\text{new}} v', C, \bar{e} \rangle \mid v := v' \} \phi'$$

which by Def. 7 is

$$\forall \bar{z}, v'. wf(G.hist) \Rightarrow loc(G, o)[\phi' \xrightarrow{v', \mathcal{H}}_{v', \mathcal{H} \cdot \langle \text{this} \xrightarrow{\text{new}} v', C, \bar{e} \rangle}]$$

Clearly this is sufficient to ensure validity of the conclusion, since the universal quantifier on v' covers the value given by $ob(C, n)$.

5.2.0.5 Proof of Theorem 1. The theorem follows by lemma 3 above and by proving that if one can prove $\vdash_C I'(\mathcal{H})$ by the class rule, then $\models_C I'(\mathcal{H})$, letting $I'(\mathcal{H})$ denote $\exists \bar{w}. I(\bar{w}, \mathcal{H})$.

Consider the rule `class`. We may assume that the premises of the rule are valid. By definition, the validity of $I'(\mathcal{H})$ is

$$\forall G, o. \text{init}_{\bar{C}} \longrightarrow^* G \wedge G.hist = \mathcal{H} \wedge o \in G.obj \wedge G[o].class = C \Rightarrow I'(\mathcal{H}/o)$$

We first prove that this holds for all C objects o in states G such that $G[o].Pr = \text{empty}$. With the given operational semantics, `Pr` is `empty` for an object o when o has finished a method, or `initC`, and it can only start a new method when `Pr` is `empty`. By lemma 1 we only need to consider states with a wellformed history. We need to show that the invariant $I(\bar{w}, \mathcal{H})$ holds after the initialization and is

maintained by every methods of class C , considering any interleaved execution according to the operational semantics. The validity of the second premise gives

$$\forall G, o. wf(G.hist) \Rightarrow \llbracket \mathcal{H} = \varepsilon \xrightarrow{o:init_C} (wf_{this}(\mathcal{H}) \Rightarrow I(\bar{w}, \mathcal{H})) \rrbracket_{G,o}$$

which by lemma 2 is the same as

$$\forall G, o. wf(G.hist) \Rightarrow \llbracket \mathcal{H} = \varepsilon \xrightarrow{o:init_C} (wf_{this}(\mathcal{H}) \Rightarrow I(\bar{w}, \mathcal{H})) \rrbracket_{G,o}$$

which by definition is

$$G \xrightarrow{o:init_C} G' \wedge loc(G, o)[\mathcal{H} = \varepsilon] \Rightarrow loc(G', o)[wf_{this}(\mathcal{H}) \Rightarrow I(\bar{w}, \mathcal{H})]$$

for all states G' with wellformed histories, and \bar{z} . This states that the class invariant $I(\bar{w}, \mathcal{H})$ holds after the initialization of class C , conditioned by local well-formedness. The condition on local well-formedness follows from the global well-formedness $wf(G'.hist)$. The condition $loc(G, o)[\mathcal{H} = \varepsilon]$ follows by induction on the length of an execution showing that no object can generate events before its initial code has started.

Similarly, the validity of the third premise gives that

$$G \xrightarrow{o:body_{C,m}} G' \wedge loc(G, o)[wf_{this}(\mathcal{H}) \wedge I(\bar{w}, \mathcal{H})] \Rightarrow loc(G', o)[wf_{this}(\mathcal{H}) \Rightarrow I(\bar{w}, \mathcal{H})]$$

for all states G, G' with wellformed histories, and all o, \bar{z} . This states that the class invariant is maintained by a method m of C under the assumption of local well-formedness. As before local well-formedness follows from global well-formedness. Thus $I(\bar{w}, \mathcal{H})$, and therefore also $I'(\mathcal{H})$, hold for all C objects o in reachable states G with empty $G[o].Pr$.

It remains to show that the invariant also holds in states G where $G[o].Pr$ is nonempty. By the first premise, we have that I' is prefix-closed with respect to the history. Thus all states in between those where $G[o].Pr$ is empty will also satisfy I' . In order to ensure I' in case of nonterminating methods (or *init*), we must consider loops and other sources of non-termination. For loops it suffices to let I' be required at the beginning of each loop iteration, which we do require in the *while* axiom. The other source of nonterminating methods is the query statement; however, here the proof rule *fetch* insists that we verify I' . Thus any proof of that method (or *init*) must establish I' at his point. By lemma 3 we have that I' is valid. We may conclude that reasoning about classes is sound.

We remark that it would be sufficient to verify I' for queries where the caller of the future equals *this* as reasoning is local, and independent of the behavior of other objects. But this would require notation for expressing the caller of a future (say $u.caller$, defined by $ft(o, n).caller = o$) in the specification (and possibly programming) language. However, the verification cost of having I' in the rule for query and in the axiom for *while*, is not great since one is obliged to prove $I(\bar{w}, \mathcal{H})$ at the end of the body.

5.3 Compositional Reasoning

The class invariant $I(\bar{w}, \mathcal{H})$ for some class C is a predicate over the fields \bar{w} , the local history \mathcal{H} , as well as the formal class parameters $\bar{c}\bar{p}$ and **this**, which are constant (read-only) variables. *History invariants* $I_C(\mathcal{H})$ for instances of C , expressed as a predicate over the local history, can be derived from the class invariant by hiding fields, i.e., $\exists \bar{w}. I(\bar{w}, \mathcal{H})$.

$$I_C(\mathcal{H}) \triangleq \exists \bar{w}. I(\bar{w}, \mathcal{H})$$

Notice that the history invariants should be prefix-closed since according to the definition in Section 4.2 C -local invariant property must be satisfied by all reachable states. Consequently, $\exists \bar{w}. I(\bar{w}, \mathcal{H})$ should be weakened if needed in order to obtain prefix-closedness. Then we assume from now on that $I_C(\mathcal{H})$ is prefix-closed.

For an instance o of C with actual parameter values \bar{e} , the *object invariant* $I_{o:C(\bar{e})}(h)$ is defined by the class invariant applied to the local projection of the history and instantiating **this** and the class parameters:

$$I_{o:C(\bar{e})}(h) \triangleq I_C(h/o)_{o,\bar{e}}^{\text{this},\bar{c}\bar{p}}$$

where I_C is a prefix-closed class invariant as above, with hidden internal state \bar{w} . We consider a composition rule for a system S of objects $o : C(\bar{e})$ together with dynamically generated objects by S . The history invariant $I_S(h)$ for such a system is then given by combining the history invariants of the composed objects:

$$I_S(h) \triangleq \text{wf}(h) \bigwedge_{(o:C(\bar{e})) \in S \cup \text{ob}(h)} I_{o:C(\bar{e})}(h)$$

where the function $\text{ob} : \text{Seq}[\text{Ev}] \rightarrow \text{Set}[\text{Obj} \times \text{Cls} \times \text{List}[\text{Data}]]$ returns the set of created objects (each given by its object identity, associated class and class parameters) in a history:

$$\begin{aligned} \text{ob}(\varepsilon) &\triangleq \{ \text{main} : \text{Main}(\varepsilon) \} \\ \text{ob}(h \cdot \langle o \xrightarrow{\text{new}} o', C, \bar{e} \rangle) &\triangleq \text{ob}(h) \cup \{ o' : C(\bar{e}) \} \\ \text{ob}(h \cdot \text{others}) &\triangleq \text{ob}(h) \end{aligned}$$

(where **others** matches all other events). By choosing S as $\{ \text{main} : \text{Main}(\varepsilon) \}$ we may reason about a global system by means of $I_{\{ \text{main} : \text{Main}(\varepsilon) \}}(h)$.

The well-formedness property serves as a connection between the local histories. Note that the system invariant is obtained directly from the history invariants of the composed objects, without any restrictions on the local reasoning, since the local histories are disjoint. This ensures compositional reasoning. The composition rule is similar to [18], which also considers dynamically created objects.

5.4 Soundness Proof of Compositional Reasoning

The proof rule for composition is:

$$\text{composition} \frac{\vdash_C I_C(h), \text{ for each } C \text{ in } \overline{Cl}}{\overline{Cl} \vdash wf(h) \bigwedge_{(o:C(\bar{e})) \in ob(h)} I_{o:C(\bar{e})}(h)}$$

Note that $\vdash_C I_C(h)$ is trivial for $I_C(h) \triangleq true$, thus one may provide invariants for a subset of the classes and using *true* as default invariant for the rest.

Theorem 2. *The object composition rule is sound.*

5.4.0.1 Proof. We show that the composition rule preserves soundness. For each class C we may then assume $\models_C I_C(h)$ which by definition is

$$\forall G, o. \text{init}_{\overline{Cl}} \longrightarrow^* G \wedge G.hist = h \wedge o \in G.obj \wedge G[o].class = C \Rightarrow I_{o:C(\bar{e})}(h)$$

Next we prove $\models I_{o:C(\bar{e})}(h)$ for all C -objects in h , i.e.,

$$\forall G. \text{init}_{\overline{Cl}} \longrightarrow^* G \wedge G.hist = h \Rightarrow \bigwedge_{(o:C(\bar{e})) \in ob_C(h)} I_{o:C(\bar{e})}(h)$$

letting $ob_C(h)$ denote the set of all C -objects in h . This reduces to proving that each C -object in $G.hist$ is found in $G.obj$. This can be proved by induction on the length of an execution. Finally by Lemma 1 we have $\models wf(h)$; and since conjunction commutes with validity we have $\models wf(h) \bigwedge_{(o:C(\bar{e})) \in ob(h)} I_{o:C(\bar{e})}(h)$.

5.5 Example

In this example we consider object systems based on the classes found in Listing 1.1 and 1.2. In order to prove that MapReduce really does output the correct number of occurrences of each word in the collection of documents, each class should guarantee the corresponding functional correctness. For example, the method *invokeMap* does take a file and emit a list of pairs such that each word in the file is associated with a counting number “1”. Moreover, we need to specify class invariants which capture the concurrent interaction between the Worker objects and the MapReduce object. For instance, the Reduce phase handled by the method *mapReduce* will start only after the Map phase has been completed. In this paper the functional correctness is given by assumption and we focus on the compositional proof based on histories such that we can derived the system property mentioned above.

Assume that the global system consists of the objects $w_1 : Worker$, $w_2 : Worker$, $w_3 : Worker$, $mr : MapReduce(wp)$, and $m : Main(mr)$, where the only visible activity of m is that it invokes *mapReduce* method on the object mr . The semantics may lead to several global histories for this system, depending on the

interleaving of the different object activities. For convenience, below we abbreviate the method names *mapReduce* to *mR*, *invokeMap* to *ivM*, and *invokeReduce* to *ivR*. One global history h caused by a call to *mR* on mr from m is as follows:

$$\begin{aligned}
& \langle m \rightarrow mr, u_1, \mathbf{mR}, \bar{e}_1 \rangle, \langle m \rightarrow mr, u_1, \mathbf{mR}, \bar{e}_1 \rangle, \\
& \langle mr \rightarrow w_1, u_2, \mathbf{ivM}, \bar{e}_2 \rangle, \langle mr \rightarrow w_2, u_3, \mathbf{ivM}, \bar{e}_3 \rangle, \\
& \langle mr \rightarrow w_2, u_3, \mathbf{ivM}, \bar{e}_3 \rangle, \langle mr \rightarrow w_1, u_2, \mathbf{ivM}, \bar{e}_2 \rangle, \\
& \langle \leftarrow w_1, u_2, \mathbf{ivM}, e_2 \rangle, \langle \leftarrow w_2, u_3, \mathbf{ivM}, e_3 \rangle, \langle mr \leftarrow, u_2, e_2 \rangle, \langle mr \leftarrow, u_3, e_3 \rangle, \\
& \langle mr \rightarrow w_2, u_4, \mathbf{ivR}, \bar{e}_4 \rangle, \langle mr \rightarrow w_2, u_4, \mathbf{ivR}, \bar{e}_4 \rangle, \langle mr \rightarrow w_1, u_5, \mathbf{ivR}, \bar{e}_5 \rangle, \\
& \langle mr \rightarrow w_3, u_6, \mathbf{ivR}, \bar{e}_6 \rangle, \langle mr \rightarrow w_3, u_6, \mathbf{ivR}, \bar{e}_6 \rangle, \langle \leftarrow w_2, u_4, \mathbf{ivR}, e_4 \rangle, \\
& \langle \leftarrow w_3, u_6, \mathbf{ivR}, e_6 \rangle, \langle mr \rightarrow w_1, u_5, \mathbf{ivR}, \bar{e}_5 \rangle, \langle \leftarrow w_1, u_5, \mathbf{ivR}, e_5 \rangle, \\
& \langle mr \leftarrow, u_4, e_4 \rangle, \langle mr \leftarrow, u_6, e_6 \rangle, \langle mr \leftarrow, u_5, e_5 \rangle, \langle \leftarrow mr, u_1, \mathbf{mR}, e_1 \rangle
\end{aligned}$$

It follows that the Reduce phase will start only after the Map phase has been completed. In addition, none of the requests sent out to the workers is uncompleted when the call to *mR* on mr is finished. We may derive these properties within the proof system from the following class invariants:

$$\begin{aligned}
I_{Worker}(\mathcal{H}) & \triangleq \mathcal{H} \leq [\langle c \rightarrow \mathbf{this}, u_1, \mathbf{ivM}, \bar{e}_1 \rangle, \langle \leftarrow \mathbf{this}, u_1, \mathbf{ivM}, e_1 \rangle] \\
& \quad \langle c \rightarrow \mathbf{this}, u_2, \mathbf{ivR}, \bar{e}_2 \rangle, \langle \leftarrow \mathbf{this}, u_2, \mathbf{ivR}, e_2 \rangle . \mathbf{some} \ c, u_1, u_2, \bar{e}_1, e_1, \bar{e}_2, e_2]^* \\
I_{MapReduce(wp)}(\mathcal{H}) & \triangleq \mathcal{H} \leq [\langle c \rightarrow \mathbf{this}, d, \mathbf{mR}, \bar{e}_1 \rangle, \langle \mathbf{this} \rightarrow _, _, \mathbf{ivM}, _ \rangle^a, \\
& \quad \langle \mathbf{this} \leftarrow, _, _ \rangle^a, \langle \mathbf{this} \rightarrow _, _, \mathbf{ivR}, _ \rangle^b, \langle \mathbf{this} \leftarrow, _, _ \rangle^b, \\
& \quad \langle \leftarrow \mathbf{this}, d, \mathbf{mR}, e_1 \rangle . \mathbf{some} \ c, d]^*
\end{aligned}$$

Here we use regular expression notation to express patterns over the history, letting $|$ denote choice, letting superscript b specify b repetitions of a pattern, and $h \leq p^*$ express that h is a prefix of a repeated pattern p where additional variables occurring in p (after **some**) may change for each repetition. Notice that the class invariant of *MapReduce* ensures that for each of the invocation event in $\langle \mathbf{this} \rightarrow _, _, \mathbf{ivM}, _ \rangle^a$, there is a corresponding fetch event in $\langle \mathbf{this} \leftarrow, _, _ \rangle^a$ by the same future identity. Same approach is applied to $\langle \mathbf{this} \rightarrow _, _, \mathbf{ivR}, _ \rangle^b$ and $\langle \mathbf{this} \leftarrow, _, _ \rangle^b$. These class invariants are straightforwardly verified in the above proof system.

The corresponding object invariants for $w_1 : Worker$, $w_2 : Worker$, $w_3 : Worker$ and $mr : MapReduce(wp)$ are obtained by substituting actual values for **this** and class parameters:

$$\begin{aligned}
I_{w_i:Worker}(h) & \triangleq h/w_i \leq [\langle _ \rightarrow w_i, u_1, \mathbf{ivM}, \bar{e}_1 \rangle, \langle \leftarrow w_i, u_1, \mathbf{ivM}, e_1 \rangle] \\
& \quad \langle _ \rightarrow w_i, u_2, \mathbf{ivR}, \bar{e}_2 \rangle, \langle \leftarrow w_i, u_2, \mathbf{ivR}, e_2 \rangle . \mathbf{some} \ u_1, u_2, \bar{e}_1, e_1, \bar{e}_2, e_2]^* \\
I_{mr:MapReduce(wp)}(h) & \triangleq h/mr \leq [\langle _ \rightarrow mr, d, \mathbf{mR}, \bar{e}_1 \rangle, \langle mr \rightarrow _, _, \mathbf{ivM}, _ \rangle^a, \\
& \quad \langle mr \leftarrow, _, _ \rangle^a, \langle mr \rightarrow _, _, \mathbf{ivR}, _ \rangle^b, \langle mr \leftarrow, _, _ \rangle^b, \\
& \quad \langle \leftarrow mr, d, \mathbf{mR}, e_1 \rangle . \mathbf{some} \ d]^*
\end{aligned}$$

The global invariant of a system S with the objects, $w_1 : Worker$, $w_2 : Worker$, $w_3 : Worker$, $mr : MapReduce(wp)$ and $m : Main(mr)$ is then

$$I_S(h) \triangleq wf(h) \wedge I_{m:Main(mr)}(h) \wedge I_{mr:MapReduce(wp)}(h) \bigwedge_{i \in \{1,2,3\}} I_{w_i:Worker}(h)$$

where well-formedness allows us to relate the different object histories. From this global invariant we may derive that the Reduce phase will start only after the Map phase has been completed. Besides, none of the requests sent out to the workers is uncompleted when the call to `mR` on `mr` is finished.

As a special case, we consider a system where the instance of `Main` invokes `mR` only once, i.e. $I_{m:Main(mr)}(h) \triangleq h/m \leq [\langle m \rightarrow mr, u, \text{mR}, \bar{e} \rangle . \mathbf{some} \ u]$. History well-formedness then ensures that the cycles defined by the remaining invariants are repeated at most once, and that variables in the patterns are connected, i.e., the future u in $I_{m:Main(mr)}$ is identical to the future d in $I_{mr:MapReduce(wp)}$. The global invariant then reduces to the following:

$$\begin{aligned}
I_S(h) \triangleq & wf(h) \wedge h/m \leq [\langle m \rightarrow mr, u, \text{mR}, \bar{e} \rangle] \wedge \\
& h/w_1 \leq [\langle mr \rightarrow w_1, u_1, \text{ivM}, \bar{e}_1 \rangle, \langle \leftarrow w_1, u_1, \text{ivM}, e_1 \rangle | \\
& \langle mr \rightarrow w_1, u_2, \text{ivR}, \bar{e}_2 \rangle, \langle \leftarrow w_1, u_2, \text{ivR}, e_2 \rangle . \mathbf{some} \ u_1, u_2, \bar{e}_1, e_1, \bar{e}_2, e_2]^* \wedge \\
& h/w_2 \leq [\langle mr \rightarrow w_2, u_3, \text{ivM}, \bar{e}_3 \rangle, \langle \leftarrow w_2, u_3, \text{ivM}, e_3 \rangle | \\
& \langle mr \rightarrow w_2, u_4, \text{ivR}, \bar{e}_4 \rangle, \langle \leftarrow w_2, u_4, \text{ivR}, e_4 \rangle . \mathbf{some} \ u_3, u_4, \bar{e}_3, e_3, \bar{e}_4, e_4]^* \wedge \\
& h/w_3 \leq [\langle mr \rightarrow w_3, u_5, \text{ivM}, \bar{e}_5 \rangle, \langle \leftarrow w_3, u_5, \text{ivM}, e_5 \rangle | \\
& \langle mr \rightarrow w_3, u_6, \text{ivR}, \bar{e}_6 \rangle, \langle \leftarrow w_3, u_6, \text{ivR}, e_6 \rangle . \mathbf{some} \ u_5, u_6, \bar{e}_5, e_5, \bar{e}_6, e_6]^* \wedge \\
& h/mr \leq [\langle m \rightarrow mr, u, \text{mR}, \bar{e} \rangle, \langle mr \rightarrow _, _, \text{ivM}, _ \rangle^a, \\
& \langle mr \leftarrow _, _, _ \rangle^a, \langle mr \rightarrow _, _, \text{ivR}, _ \rangle^b, \langle mr \leftarrow _, _, _ \rangle^b, \\
& \langle \leftarrow mr, u, \text{mR}, e \rangle]
\end{aligned}$$

This invariant allows a number of global histories, depending on the interleaving of the activities in the different objects. The history h presented first in this section satisfies the invariant, and represents one particular interleaving.

Based on the assumption of functional correctness of each class, we now can derive that the MapReduce object does output the correct number of occurrences of each word in the collection of documents.

6 Related Work

Models for asynchronous communication without futures have been explored for process calculi with buffered channels [25], for agents with message-based communication [2], for method-based communication [?], and in particular for Java [22]. Behavioral reasoning about distributed and object-oriented systems is challenging, due to the combination of concurrency, compositionality, and object orientation. Moreover, the gap in reasoning complexity between sequential and distributed, object-oriented systems makes tool-based verification difficult in practice. A survey of these challenges can be found in [4]. Soundness of the parallel composition rules for shared-variable concurrency and synchronous message passing are proved in [16]. A Hoare Logic for concurrent processes (objects) is presented in [14]. The Hoare Logic is compositional, and soundness and relative completeness are proven. In contrast to our work, communication is by message passing rather than by futures, and the objects communicate through FIFO channels.

The present approach follows the line of work based on communication histories to model object communication events in a distributed setting [8, 10, 25]. Objects are concurrent and interact solely by method calls and futures, and remote access to object fields are forbidden. By creating unique references for method calls, the *label* construct of Creol [29] resembles futures, as callers may postpone reading result values. Verification systems capturing Creol labels can be found in [4, 21]. However, a label reference is local to the caller, and cannot be shared with other objects. A reasoning system for futures has been presented in [15], using a combination of global and local invariants. Futures are treated as visible objects rather than reflected by events in histories. In contrast to our work, global reasoning is obtained by means of global invariants, and not by compositional rules. Thus the environment of a class must be known at verification time. SCOOP [?, ?] and Cameo [?] are two concurrency models for Eiffel [?] based on the concepts of design-by-contract. Compared with our work, these two approaches are not using histories.

A reasoning system for asynchronous methods in ABS without futures is presented in [18]. We here define a five-event semantics reflected actions on shared futures and object creation. The semantics gives a clean separation of the activities of the different objects, which leads to disjointness of local histories. Thus, object behavior can be specified in terms of the observable interaction of the current object only. This is essential for obtaining a simple reasoning system. In related approaches, e.g., [4, 21], events are visible to more than one object. The local histories must then be updated with the activity of other objects, resulting in more complex reasoning systems. Based on the five-event semantics, we present a compositional reasoning system for distributed, concurrent objects with asynchronous method calls. A class invariant defines a relation between the inner state and the observable communication of instances, and can be verified independently for each class. The class invariant can be instantiated for each object of the class, resulting in a history invariant over the observable behavior of the object. Compositional reasoning is ensured as history invariants may be combined to form global system specifications. The composition rule is similar to [18], which is inspired by previous approaches [33, 34]. This work is an extension of our former paper [20]. Here we analyze a larger case study using futures, extend the language and semantics, including object creation, branching and looping constructs. Also, soundness proofs for class reasoning and in particular object composition are provided.

7 Conclusion

In this paper we have considered concurrent objects communicating by means of futures and a notion of non-blocking methods calls. This concurrency model is different from that found in mainstream languages such as Java. We find it interesting since it is based on high-level synchronization primitives, rather than locks and signaling, and it allows the caller to control the waiting time by means of different ways of calling a method, suspending, blocking, or non-blocking. In

addition it directly supports distribution, autonomy, message-based communication, and object-orientation. Thus the class mechanism is devoted to programming of concurrent and autonomous objects, whereas internal data structures are programmed by the use of data types. This concurrency model has recently been the theme of several EU projects, including Credo, Hats, Envisage, and Upscale. Tool support for this concurrency model have been investigated in several ways, including compilation of the ABS language to more low-level languages including Java, Erlang, and Scala. The concurrency model gives rise to a clean, compositional semantics. Compositionality is a key property for scalability, allowing program units to be developed, tested, and understood, independently. The concurrency model may be relevant for Java, extended with asynchronous methods, and restricted so that remote access, notification, and explicit locking, are avoided.

The focus of this paper is program reasoning, and the concurrency model is chosen due to advantages with respect to program reasoning, while supporting true concurrency of objects. Compositional reasoning is facilitated by expressing object properties in terms of observable interaction between the object and its environment, recorded on communication histories. Object generation is reflected in the history by means of object creation events. A method call cycle with multiple future readings is reflected by four kinds of events, giving rise to disjoint communication alphabets for different objects. Specifications in terms of history invariants may then be derived independently for each object and composed in order to derive properties for concurrent object systems. At the class level, invariants define relationships between class attributes and the observable communication of class instances. The presented reasoning system is proved sound with respect to the given operational semantics. This system is easy to apply in the sense that class reasoning is similar to standard sequential reasoning, but with the addition of effects on the local history for statements involving futures. In particular, reasoning inside classes is not affected by the complexity of concurrency and synchronization; and one may express assumptions about inputs from the environment when convenient.

The main result of this paper is soundness of the rule for composition objects running in parallel. A minor result is that sound composition requires the `query` rule to have a condition related to the invariant, not found in earlier papers. We consider here global history invariants that are continuously satisfied, in the sense that any reachable global configuration of an object system must satisfy the invariant. The condition on the `query` rule would not be needed with a weaker notion of global history invariants stating that the global invariant holds as long as all objects are live (not blocked). Verification-wise the condition on the `query` rule, is somewhat similar to a query statement releasing the processor, as for instance the `await future` statement of the ABS language. Semantically, a blocking query has the advantage that it does not change the state, whereas a non-blocking query gives a state satisfying the local invariant. Thus the combination of the query and processor release mechanisms will not add significant verification complexity, and is also attractive from a programming perspective.

In order to focus on the future mechanism, this paper considers a core language with shared futures. The report version [19] considers a richer language, including constructs for inter-object process control and processor release. The verification system is suitable for an implementation within the KeY framework. With support for (semi-)automatic verification, such an implementation will be valuable when developing larger case studies. It is also natural to investigate how our reasoning system would benefit from extending it with rely/guarantee style reasoning [16]. Assumptions about callee behavior may, for instance, be used to express properties of return values. More sophisticated techniques may also be used, e.g., [13, 28] adapts rely/guarantee style reasoning to history invariants. However, such techniques require more complex composition rules.

References

1. E. Ábrahám, I. Grabe, A. Grüner, and M. Steffen. Behavioral interface description of an object-oriented language with futures and promises. *Journal of Logic and Algebraic Programming*, 78(7):491–518, 2009.
2. G. Agha, S. Frolund, W. Kim, R. Panwar, A. Patterson, and D. Sturman. Abstraction and modularity mechanisms for concurrent computing. *Parallel Distributed Technology: Systems Applications, IEEE*, 1(2):3–14, may 1993.
3. A. Ahern and N. Yoshida. Formalising java rmi with explicit code mobility. *Theoretical Computer Science*, 389(3):341–410, 2007. Semantic and Logical Foundations of Global Computing.
4. W. Ahrendt and M. Dylla. A system for compositional verification of asynchronous objects. *Science of Computer Programming*, 2010.
5. B. Alpern and F. B. Schneider. Defining liveness. *Information Processing Letters*, 21(4):181–185, Oct. 1985.
6. H. G. Baker Jr. and C. Hewitt. The incremental garbage collection of processes. In *Proceedings of the 1977 symposium on Artificial intelligence and programming languages*, pages 55–59, New York, NY, USA, 1977. ACM.
7. B. Beckert, R. Hähnle, and P. H. Schmitt, editors. *Verification of Object-Oriented Software: The KeY Approach*, volume 4334 of *LNCS*. Springer, 2007.
8. M. Broy and K. Stølen. *Specification and Development of Interactive Systems*. Monographs in Computer Science. Springer, 2001.
9. M. Clavel, F. Durán, S. Eker, P. Lincoln, N. Martí-Oliet, J. Meseguer, and C. L. Talcott. *All About Maude - A High-Performance Logical Framework, How to Specify, Program and Verify Systems in Rewriting Logic*, volume 4350 of *LNCS*. Springer, 2007.
10. O.-J. Dahl. Can program proving be made practical? In M. Amirchahy and D. Néel, editors, *Les Fondements de la Programmation*, pages 57–114. Institut de Recherche d’Informatique et d’Automatique, Toulouse, France, Dec. 1977.
11. O.-J. Dahl. Object-oriented specifications. In *Research directions in object-oriented programming*, pages 561–576. MIT Press, Cambridge, MA, USA, 1987.
12. O.-J. Dahl. *Verifiable Programming*. International Series in Computer Science. Prentice Hall, New York, N.Y., 1992.
13. O.-J. Dahl and O. Owe. Formal methods and the RM-ODP. Research Report 261, Department of Informatics, University of Oslo, Norway, May 1998.
14. F. S. de Boer. A Hoare Logic for Dynamic Networks of Asynchronously Communicating Deterministic Processes. *Theoretical Computer Science*, 274:3–41, 2002.

15. F. S. de Boer, D. Clarke, and E. B. Johnsen. A complete guide to the future. In R. de Nicola, editor, *Proc. 16th European Symposium on Programming (ESOP'07)*, volume 4421 of *LNCS*, pages 316–330. Springer, Mar. 2007.
16. W.-P. de Roever, F. de Boer, U. Hannemann, J. Hooman, Y. Lakhnech, M. Poel, and J. Zwiers. *Concurrency verification: introduction to compositional and non-compositional methods*. Cambridge University Press, New York, NY, USA, 2001.
17. J. Dean and S. Ghemawat. Mapreduce: simplified data processing on large clusters. *Commun. ACM*, 51(1):107–113, Jan. 2008.
18. C. C. Din, J. Dovland, E. B. Johnsen, and O. Owe. Observable behavior of distributed systems: Component reasoning for concurrent objects. *Journal of Logic and Algebraic Programming*, 81(3):227–256, 2012.
19. C. C. Din, J. Dovland, and O. Owe. An approach to compositional reasoning about concurrent objects and futures. Research Report 415, Dept. of Informatics, University of Oslo, Feb. 2012. Available at <http://folk.uio.no/crystal/>.
20. C. C. Din, J. Dovland, and O. Owe. Compositional reasoning about shared futures. In G. Eleftherakis, M. Hinchey, and M. Holcombe, editors, *Proc. International Conference on Software Engineering and Formal Methods (SEFM'12)*, volume 7504 of *LNCS*, pages 94–108. Springer, 2012.
21. J. Dovland, E. B. Johnsen, and O. Owe. Verification of concurrent objects with asynchronous method calls. In *Proceedings of the IEEE International Conference on Software Science, Technology & Engineering (SwSTE'05)*, pages 141–150. IEEE Computer Society Press, Feb. 2005.
22. K. E. K. Falkner, P. D. Coddington, and M. J. Oudshoorn. Implementing asynchronous remote method invocation in java, 1999.
23. R. H. Halstead Jr. Multilisp: a language for concurrent symbolic computation. *ACM Trans. Program. Lang. Syst.*, 7(4):501–538, Oct. 1985.
24. Full ABS Modeling Framework (Mar 2011). Deliverable 1.2 of project FP7-231620 (HATS), available at <http://www.hats-project.eu>.
25. C. A. R. Hoare. *Communicating Sequential Processes*. International Series in Computer Science. Prentice Hall, 1985.
26. International Telecommunication Union. Open Distributed Processing - Reference Model parts 1–4. Technical report, ISO/IEC, Geneva, July 1995.
27. A. S. A. Jeffrey and J. Rathke. Java Jr.: Fully abstract trace semantics for a core Java language. In *Proc. European Symposium on Programming*, volume 3444 of *LNCS*, pages 423–438. Springer, 2005.
28. E. B. Johnsen and O. Owe. Object-oriented specification and open distributed systems. In O. Owe, S. Krogdahl, and T. Lyche, editors, *From Object-Orientation to Formal Methods: Essays in Memory of Ole-Johan Dahl*, volume 2635 of *LNCS*, pages 137–164. Springer, 2004.
29. E. B. Johnsen and O. Owe. An asynchronous communication model for distributed concurrent objects. *Software and Systems Modeling*, 6(1):35–58, Mar. 2007.
30. B. H. Liskov and L. Shriram. Promises: Linguistic support for efficient asynchronous procedure calls in distributed systems. In D. S. Wise, editor, *Proc. SIGPLAN Conference on Programming Language Design and Implementation (PLDI'88)*, pages 260–267. ACM Press, June 1988.
31. J. Meseguer. Conditional rewriting logic as a unified model of concurrency. *Theoretical Computer Science*, 96:73–155, 1992.
32. B. Morandi, S. S. Bauer, and B. Meyer. Scoop - a contract-based concurrent object-oriented programming model. In P. MÅ¼ller, editor, *Advanced Lectures on Software Engineering, LASER Summer School 2007/2008*, volume 6029 of *Lecture Notes in Computer Science*, pages 41–90. Springer, 2008.

33. N. Soundararajan. Axiomatic semantics of communicating sequential processes. *ACM Transactions on Programming Languages and Systems*, 6(4):647–662, Oct. 1984.
34. N. Soundararajan. A proof technique for parallel programs. *Theoretical Computer Science*, 31(1-2):13–29, May 1984.
35. A. Yonezawa, J.-P. Briot, and E. Shibayama. Object-oriented concurrent programming in ABCL/1. In *Conference on Object-Oriented Programming Systems, Languages and Applications (OOPSLA '86)*. *Sigplan Notices*, 21(11):258–268, Nov. 1986.