



ELSEVIER

Contents lists available at ScienceDirect

Science of Computer Programming

www.elsevier.com/locate/scico


A formal actor-based model for streaming the future

Keyvan Azadbakht^{a,b,*}, Frank S. de Boer^a, Nikolaos Bezirgiannis^a,
Erik de Vink^c

^a Centrum Wiskunde & Informatica, Amsterdam, the Netherlands

^b Leiden University, Leiden, the Netherlands

^c Eindhoven University of Technology, Eindhoven, the Netherlands



ARTICLE INFO

Article history:

Received 7 January 2019

Received in revised form 14 October 2019

Accepted 21 October 2019

Available online xxxx

Keywords:

Future

Streaming

Cooperative scheduling

Active objects

Social networks

ABSTRACT

Asynchronous Actor-based programming has gained increasing attention as a model of concurrency and distribution. The Abstract Behavioral Specification (ABS) language is an actor-based programming language that has been developed for both the modeling and formal analysis of distributed systems. In ABS, actors are modeled as concurrent objects that communicate by asynchronous method calls. Return values are also communicated asynchronously via return statements and so-called futures.

Many modern distributed software applications require a form of continuous interaction between their components which consists of streaming data from a server to its clients. In this paper, we extend the ABS language in order to support the streaming of data. We introduce the notion of “future-based data streams” by augmenting the syntax, type system, and operational semantics of ABS. As a proof of concept, we further discuss a prototype implementation for supporting future-based data streams on top of ABS, and discuss the impact of the use of these data streams in ABS on the performance in the implementation of a distributed application for the generation of social networks.

© 2019 Elsevier B.V. All rights reserved.

1. Introduction

Since the rapid growth in big data, data streaming is widely used in many distributed applications, e.g., telecommunications, event-monitoring and detection, and sensor networks. Data streaming is a client/server pattern which, in essence, consists of a continuous generation of data by the server and a sequential and incremental processing of the data by the client. Data streams are naturally processed differently from batch data. Functions cannot operate on data streams as a whole, as the produced data can be unlimited. Hence, new programming abstractions are required for the continuous generation and consumption of data in the streams.

Data streaming is highly relevant in modern distributed systems. Actor-based languages are specifically designed for describing such systems [1]. They provide an event-driven model of concurrency where messages are communicated asynchronously and processed by pattern matching mechanism [2]. Concurrent objects generalize this model to *programming to interface* discipline by modeling messages as asynchronous method invocations. The main contribution of this paper is to integrate data streaming mechanism with concurrent object systems.

* Corresponding author at: Centrum Wiskunde & Informatica, Amsterdam, the Netherlands.

E-mail addresses: k.azadbakht@cwi.nl (K. Azadbakht), frank.s.de.boer@cwi.nl (F.S. de Boer), n.bezirgiannis@cwi.nl (N. Bezirgiannis), evink@win.tue.nl (E. de Vink).

The Abstract Behavioral Specification (ABS) language [3] has been developed for formal modeling and analysis of distributed systems. ABS is an *active object* language, where active objects communicate via asynchronous method calls, and as such support a programming to interfaces discipline, in contrast to, e.g., Erlang [2] actors where messages are untyped. In ABS, concurrent objects represent processes that execute in parallel and interact via asynchronous communication of messages. A message specifies one of the methods of the called object. Return values are also communicated asynchronously via so-called futures [4]. Futures are dynamically generated references for checking availability and reading of the return value. ABS is statically typed and supports a *programming to interface* discipline. In contrast to the *run-to-completion* mode of method execution (e.g., in [5]), ABS further provides the powerful feature of *cooperative scheduling* which allows an object to suspend in a controlled manner the current executing method invocation (also known as process) and schedule another invocation of one of its methods.

In this paper, we extend the ABS language in order to support the streaming of data between a server and its clients. We introduce “future-based data streams” which integrates futures and data streams, and which specifies the return type of so-called *streaming* methods. Upon invocation of such a method a new future is created which holds a reference to the generated stream of data. Data items are added to the stream by the execution of a *yield* statement. Such a statement takes as parameter an expression the value of which is added to the stream, *without* terminating the execution of the method. The *return* statement terminates the execution of a streaming method, and is used to signal the end of data streaming. Even though no new data is produced, the existing data values in the stream buffer can be retrieved by the consumers.

The values generated by the server (the streaming method) can be obtained incrementally and sequentially by a client by querying the future corresponding to this method invocation. By the nature of data streaming, it is natural to restrict the streaming to the asynchronous method calls. Therefore there is no support for synchronous invocation of streaming methods.

In this paper, we introduce two different implementations of streams tailored to different forms of parallel processing of data streams. Obtaining data from a *destructive* stream involves the removal of the data, whereas in a *non-destructive* stream the data persists. Which of the implementation is used is determined by the caller of the streaming method (the *creator* of the stream) which is not necessarily the consumer of the data stream. The creator can then provide the consumers with a reference to the stream. Both the streaming method (producer) and the consumers which hold a reference to the data stream are *not* exposed to the underlying implementation of the stream, i.e., these different implementations are not represented by different *types* of data streams. This allows for a separation of concerns between the generation and processing of data streams, on the one hand, and their orchestration, on the other hand. This also enables reusability of the implementation of producers and consumers for both consumption approaches.

A preliminary discussion of the overall idea underlying this paper is given in [6]. As an extension, in this paper we introduce the different implementations of data streams, an operational semantics for both implementations of streams, a new type system which formalizes the integration of futures and data streams, and a proof of type-safety. Further, we show how the basic mechanism in ABS of cooperative scheduling of asynchronously generated method invocations itself can be used to implement data streams and the cooperative scheduling of streaming methods.

As a proof of concept, exploiting a prototype implementation for supporting future-based data streams on top of ABS, we present the impact of the above-mentioned feature on the performance in the implementation of a distributed application for the generation of social networks. The notion of data streaming abstracts from the specific implementation of ABS. In our case, we make use of the distributed Haskell backend of ABS [7] for the case study on future-based data streams reported in this paper.

The overall contribution of this paper is a formal model of streaming data in the ABS language, which fully complies and generalizes the asynchronous model of computation underlying the ABS language. Since ABS is defined in terms of a formal operational semantics which supports a variety of formal analysis techniques (e.g., deadlock detection [8] and [9]), we thus obtain a general formal framework for the modeling and analysis of different programming techniques for processing data streams, e.g., *map-reduce* and *publish-subscribe* [10]. To the best of our knowledge, our work provides a first formal type system and operational semantics for a general notion of streaming data in a high-level actor-based programming language.

Plan of the Paper. This paper is organized as follows: a brief description of the ABS programming language is given in Section 2. The notion of a future-based data stream is specified as an extension of ABS in Section 3. In section 4, it is shown that the well-typedness of a program in the extended ABS is preserved. Section 5 discusses the usage of streams in a distributed setting. In section 6, an implementation of data streams is given as an API written in ABS. In Section 7, a case study on social network simulation is discussed, which uses the proposed notion of streams. Related works are discussed in section 8. Finally we conclude in section 9.

2. The ABS programming language

Here we briefly highlight the main features of ABS relevant to our work in this paper (for a more detailed description we refer to [3]). In ABS, parallel (or concurrent) processes are generated by asynchronous method calls of the form $\mathbf{Fut}\langle T \rangle \bar{f} = \circ ! m(\bar{e})$, where \bar{f} is a future used as a reference to the return value of the asynchronous call of the method m , \circ is an expression denoting the called object, \bar{e} are the actual parameters, and $\mathbf{Fut}\langle T \rangle$ is the type of futures which hold return values of type T . Such a call generates a process for the execution of the invoked method which is stored in the

$$s ::= s; s \mid x = [\mathbf{nd}] e!m(\bar{e}) \mid \mathbf{yield} e \mid \mathbf{return} \mid \\ \mathbf{await} e? \mathbf{finished} \{s\} \mid x = e.\mathbf{get} \mathbf{finished} \{s\}$$

Fig. 1. Syntax.

process pool of the called object. Concurrent objects in ABS are executing at most one process at a time. The executing process of an object is executing in parallel with all the executing processes of the other objects. An active object is created by the expression **new**C, where C denotes a class.

Further, ABS features synchronous method calls and a high-level synchronization mechanism (i.e., cooperative scheduling) which allows an object to suspend the execution of the current process, and schedule another (enabled) process for execution, by means of **await** and **suspend** statements. The **await** $f?$ statements suspends the current process, in case the future f is not resolved (yet). The **await** b statement similarly suspends the current process, in case the Boolean guard b evaluates to **False**. The suspended processes can be (re-)activated, in case f is resolved, and b evaluates to **True**, respectively. The **suspend** statement suspends the current process unconditionally. The statement $x = f.\mathbf{get}$ blocks the object till the future f has been resolved, and assigns the computed return value to the variable x .

3. Future-based data streams

In this section, we define future-based data streaming in the context of the ABS language. A *streaming* method is statically typed, namely, the return type of the method is followed by the keyword **stream**, specifying that the method streams values of that type. As mentioned before, ABS features a programming to interfaces discipline. Therefore the caller can asynchronously call a streaming method, provided that the interface of the callee includes the method definition.

Data streaming is defined as a stream of return values from a callee to the data consumers of the stream in an asynchronous fashion. An invocation of a streaming method creates a stream. The callee first create an empty stream, and then produces and stores data to the stream buffer via the **yield** statement. The caller assigns the invocation to a variable of type **Stream**<T> for the return type T **stream** of the callee. The stream variable can be passed around. Therefore different variables in multiple processes (a process is the execution of an asynchronous method call) may refer to the same stream and retrieve data from it.

We distinguish between two different kinds of streams: *destructive* and *non-destructive* streams. The kind of stream is determined by the caller upon the invocation of the streaming method. In destructive streams, values are retrieved from a FIFO queue which stores the data produced but not yet consumed. Querying availability of data values in an empty queue gives rise to a cooperative release of control (further discussed below). Also an attempt to take a value from a stream where the callee is terminated (and thus no further data streaming will take place) gives rise to the execution of a block of statements specified by programmer for this reason, thus avoiding the generation of a corresponding error (see below). Parallel processes which have access to the same destructive stream compete for the different data items produced. Consequently, the parallel processing of destructive data streams gives rise to *race conditions*, in the sense that different order of requests to read from a stream may correspondingly give rise to different data values. Note that at most one process can destructively read a specific data value. On the other hand, a non-destructive stream allows complete *sharing* of all the data produced which are only *read* to be processed. As described in more detail below, non-destructive streams maintain access by means of cursors at different positions of the buffer which allows for its asynchronous parallel processing.

Abstracting from the typing information, to be discussed in more detail below, the syntax of our proposed extension of ABS, i.e., that of future-based data streams, is specified in Fig. 1, where e denotes an expression (i.e., a variable name, etc), \bar{e} denotes a sequence of expressions, x is a variable name, m is a method name, and s denotes a sequence of statements.

In the asynchronous invocation $x = [\mathbf{nd}] e!m(\bar{e})$ of a streaming method, the optional keyword **nd** indicates the creation of a new non-destructive stream.

Execution of the **yield** statement, which can only be used in the body of a streaming method, consist of queuing the value of the specified expression.

Execution of the **return** statement by a streaming method indicates termination of the data generation which is signaled to the consumers of the stream by queuing the special value η .

The **await-finished** statement allows to check the buffer of the stream denoted by the expression e in the following manner: if there is at least one proper value, different from the signal η , in the buffer, the statement is skipped. In case the buffer is empty, the current process *suspends* such that the object can activate another process. The statement s is executed in case the buffer only contains the signal η .

The **get-finished** statement allows to actually retrieve (in case of a destructive stream) or read (in case of a non-destructive stream) a next data value. It however *blocks* the whole object when the buffer is empty. As above, statement s is executed when the buffer only contains the signal η .

In **await-finished** and **get-finished**, the keyword **finished** and its following block can be omitted if the block is empty.

We next illustrate the difference in the behaviour of destructive and non-destructive access to a stream by the following simple toy example. Consider the streaming method $m()$:

```
Int stream m() {
  yield 1; yield 2; return;
}
```

This method adds 1 and 2, followed by a termination token to the resulting stream buffer.

The following snippet asynchronously calls the above method definition `m()` on some object `o` which gives rise to two references `r1` and `r2` to the resulting destructive stream.

```
Stream<Int> r1 = o!m();
Stream<Int> r2 = r1;
```

The following code uses the above references, with the assumption that it is the only process that consumes data items of the stream:

```
Int x, y, z;
(1) Int x = r1.get finished {x = -1};
(2) Int y = r2.get finished {y = -1};
(3) Int z = r1.get finished {z = -1};
```

Once the process corresponding to the method call `m()` on (possibly remote) object `o` is executed and the results are provided to the stream, the values 1, 2, and -1 are assigned to `x`, `y` and `z`, respectively. These values are consumed from the stream and assigned to the variables incrementally as soon as they are provided by `m()`. In the above code, the object possibly blocks on any of the three statements, if a value (whether an integer or the terminating token) is not yielded to the stream yet. The statement (1) destructively reads 1 from the stream via `r1` and assigns it to `x`. The statement (2) destructively reads 2 from the same stream via the other reference `r2` and assigns it to `y`. However, the statement (3) runs the `finished` statement which assigns -1 to `z`, since it reads the terminating token (i.e., the stream is already terminated). Any further `get` operations on every variable referring to the stream also read the terminating token.

To show how a non-destructive stream works in the same setting, suppose we use the following references `r1` and `r2` in the above code (note that the keyword `nd` denotes that the resulting stream is non-destructive).

```
Stream<Int> r1 = nd o!m();
Stream<Int> r2 = r1;
```

With the same incremental production of values and blocking mechanism, in this setting the values 1, 1, and 2 are assigned to `x`, `y` and `z`, respectively. The statement (1) non-destructively reads 1 from the stream via `r1` and assigns it to `x`. The statement (2) non-destructively reads 1 from the same stream via the other reference `r2` (with its own cursor to the stream) and assigns it to `y`. Finally, the statement (3) assigns 2 to `z`, since the cursor of `r1` is already moved forward by statement (1). Note that any number of further `get` operations on `r1` will read the terminating token. It is important to observe the role of cursors per each stream variable that gives rise to such behaviour.

Note that the assignment of a non-destructive reference (`r2 = r1`) is different from the standard ABS assignment in the sense that, in addition to the assignment of the reference to stream, it also assigns the cursor. Based on this design, the cursor copying is required as each stream variable represents a new access to the stream to *all* data values from the position its cursor denotes.

3.1. Design decisions

Integration of streams with ABS, where we enjoy the advantages of both, roots in the ever-growing application of data streaming in different domains. The consumption approaches of the stream (i.e., destructive or non-destructive) are not fixed in the streams in form of different data types. Instead, the creator of the stream determines the consumption approach of the stream instance, in order to maintain generality. Note that the creator of the stream is not necessarily the consumer of the stream, and by design, it can be considered as part of the producer process (e.g., using *factory method* design pattern) that forces one of the above consumption approaches to the consumers.

We support both destructive and non-destructive data streams, as they can be naturally used to implement, respectively, *one-of-n* semantics (only one consumer reads a given data as in, e.g., *data parallelism* model), and *one-to-n* message delivery (a given data can be read by all such consumers as in, e.g., *one-to-many trainer and learners* and *publish/subscribe* model). Also integration of data streaming and cooperative scheduling enables enhancing concurrency and parallelism on the consumer side.

Note that the above two approaches allow for designing a third hybrid consumption approach where, at the intra-object level, every access to the stream buffer is via an object field (shared variable), and at the inter-object level, the cursor is copied (i.e., via passing parameters in method invocations).

```

interface IMapper<T> {
  Int stream map(Stream<T> s);
}

interface IReducer<T> {
  Pair<Int, Int> reduce(Stream<T> s);
}

interface IPar<T> {
  Int start(Stream<T> s, Int num);
}

interface IProd<T> {
  T stream streamer();
}

class Mapper()
implements IMapper<T> {
  Int stream map(Stream<T> s) {
    Bool last = False;
    while(last==False){
      T v = s.get finished {last=True};
      if (last == False) {
        yield v.value();
      }
    }
    return;
  }
}

class Reducer()
implements IReducer<T> {
  Pair<Int, Int> reduce(Stream<Int> s)
  {
    Bool last = False;
    Int count = 0;
    while(last==False){
      Int v = s.get
        finished {last=True};
      if (last == False){
        count = count + 1;
        sum = sum + v;
      }
    }
    return Pair(sum, count);
  }
}

class Producer implements IProd<T>
{
  T stream streamer() {
    // yields a seq of data of type T
    return;
  }
}

class Par implements IPar<T> {
  Int start(Stream<T> s, Int num) {
    Int m = 1;
    Int sum = 0, avg = 0, count = 0;
    List<Fut<Pair<Int, Int>>> l = Nil;
    while(m<=num) {
      IMapper<T> p = new Mapper();
      Stream<Int> s2 = p!map(s);
      IReducer<T> q = new Reducer();
      l = Cons(q ! reduce(s2), l);
      m=m+1;
    }
    while (l != Nil) {
      Pair<Int,Int> pair = head(l).get;
      case (pair) {
        Pair(a, b) => {
          sum = sum + a;
          count = count + b;
        }
      }
      l = tail(l);
    }
    if (count > 0) return sum / count;
    else return 0;
  }
}

{// Main block
  IProd<T> producer = new Producer();
  Stream<T> s = producer ! streamer();
  IPar<T> par = new Par();
  Int average = par.start(s, 4);
}

```

Fig. 2. Parallel data processing based on Map-Reduce data model.

3.2. Example of destructive streams

The code example in Fig. 2 illustrates the use of ABS destructive data streams in modeling a parallel *map-reduce* processing of a data stream. The mapping step maps each streamed data value of type *T* to a data value of type *Int*, and the reduction step calculates the average of those integers.

An ABS program is a set of interface and class definitions, followed by the main block of the program, which is an anonymous block at the end of the program. The main block is the initial run-time process (similar to `public static void main` in java).

Each class implements at least one interface. The type of a reference variable to an object can only be an interface, and the object must be an instance of a class that implements the interface. Every object instance is an active object, namely, it features a dedicated thread of control, and can have (at most) one active process among its processes. Each process of an object is initiated by an asynchronous call of a method.

The program is composed of four interfaces: *IProd* types a class with a streaming method to stream the data values of type *T* to be processed. The interface *IPar* types a class for spawning multiple chains of active objects for map-reduce pro-

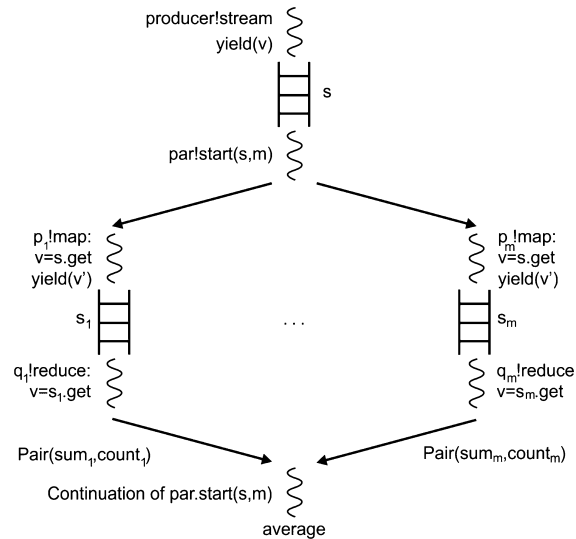


Fig. 3. Control and data flow.

processing. Each chain is a pipeline processing of the data values retrieved from the stream which is shared among the chains. The interfaces `IMapper` and `IReducer` type the objects that form a pipeline chain. These interfaces are implemented by four classes `Producer`, `Par`, `Mapper` and `Reducer`, respectively. The above definitions are followed by the main block of the program. As shown in the main block, the general idea is that the data values of the stream `s` will be processed in parallel by `num` computationally identical pipelines, and the aggregated result, which is the average of those values, is returned as the final result. Runtime control and data flow of the example are also illustrated in Fig. 3, where each thread represents a process created by an asynchronous method call.

The asynchronous invocation of method `streamer` on `producer` in the main block returns a reference `s` to a stream. The method `start` provided by `IPar<T>` enables parallel processing of the stream by creating multiple chains (`num`) of two active objects of type `IMapper<T>` and `IReducer<T>`, where the former retrieves values from `s`, and yields a mapped value of type integer to an intermediate stream `s2`, and the latter consumes those integers whose sum and count are reduced to a pair which is the sum and the count of those integers processed by one chain. The futures of the pairs resulting from calling `reduce` in different chains are stored in list `l`. The elements of the list are then used as a synchronization means, namely, awaiting until each process resolves the corresponding future by providing the return value. Finally the average is calculated by the `start` from the reduced pairs.

Similar to parallel map-reduce transformations on streams in languages like Scala, the following pseudo-code can be used as a simplified abstract replacement for the code in 2:

```
s.par(num).map(_.value).average();
```

where a sequence of *transformation* methods (e.g., `map` and `filter`) followed by a *reduce* method (e.g., `count` and `average`) can be executed in parallel by `num` threads (modeled by active objects) on stream `s`.

Note that our implementation utilizes two ways of parallelism: 1) *horizontal parallelism*, which is achieved by creating multiple chains of active objects, e.g., `pi` and `qi` and intermediate streams `si` in Fig. 3, and 2) *vertical parallelism*, which is achieved by pipeline processing, e.g., the process `map` in `pi` that yields values to `si` runs in parallel with `reduce` in `qi` that consumes the values immediately upon their availability.

3.3. Example of non-destructive streams

In the example specified in Fig. 4, we represent a basic means of *publish/subscribe* communication via non-destructive streams in a social network such as Twitter. An object of class `Member` denotes a member in the network that can follow and be followed by multiple members. The main idea is to implement each member object such that: 1) it can follow multiple members by reading their stream of posts 2) its stream of posts can be read by multiple members that follow the member 3) it can post new items to its stream. The object naturally needs to interleave these tasks. To this aim, each member is modeled as an actor with a process to post new items to its stream (`share`), a set of processes one per each member it follows, in order to read their streams (`follow`), and a set of processes from other members that request to follow the member (`request`). These processes can be interleaved by the *ABS cooperative scheduling*. The active process can cooperatively release control conditionally, e.g., the `await` statement in `follow` which checks whether there is no new post to be read from a specific member, or unconditionally, e.g., the `suspend` in `share` after posting a new item gives rise to unconditional release of control. In both cases, other processes of the member object can be activated.

```

interface IMem {
  Unit run();
  Unit follow(IMem p);
  Stream<Post> request();
}

class Member implements IMem
{
  Stream<Post> r;
  // r is a stream of
  // posts for followers

  Unit run() {
    if (r == null)
      r = nd this!share();
  }

  Post stream share() {
    Post post;
    while(True) {
      // Next post is ready
      yield post;
      suspend;
    }
    return;
  }
}

Unit follow(IMem p) {
  Fut<Stream<Post>> f =
  p ! request();
  await f?;
  Stream<Post> r = f.get;
  Bool last = False;
  while(last = False) {
    await r? finished
    {
      //probably p left!
      last = true;
    }
    Post post = r.get;
    // consume post
  }
}

Stream<Post> request() {
  // accept as a follower
  return r;
}

```

Fig. 4. Parallel data processing based on publish/subscribe pattern.

The method `follow` sends a request to a member denoted by the argument p . The data (i.e., posts) can be retrieved from the resulting stream r of the member p by the current member. In other words, the current member object *follows* object p . Further, a followed member returns a reference to the same data stream for all the followers, denoted by r in the class `Member`. Each follower uses its corresponding cursor to read from the stream belonging to the followed member. Note the difference between the return types of `share` and `request`. The former is a streaming method that creates and populates a stream, and can *only* be called asynchronously with the return type `Stream<Post>`, whereas the latter is a non-streaming method that returns a reference to an existing stream, and returns `Stream<Post>` or `Fut<Stream<Post>>`, respectively, depending on being called synchronously or asynchronously.

By `await` statement, a *follow* process queries the availability of the next post that is *new* from the perspective of the non-destructive stream variable r , denoted by the variable cursor. If the new post is available it is retrieved and consumed. Otherwise the process is suspended so that another enabled process is activated. As such, the member receives posts from all the members it follows, processes the follow *requests* of other members, and posts new data. The stream corresponding to a followed member can signal the termination. In such case, the `follow` process in the follower object which corresponds to a followed member terminates after retrieving the remaining posts, as the *finished* block of the `await` statement falsifies the loop condition. The process that instantiates a new member (not mentioned here) also initiates the member by calling the `run` which itself calls the `share` method which returns a new stream and continuously adds new posts to it.

3.4. Type system

The ABS type system is presented in [3]. An extension of the type system is specified below using the same notation, which types the streams and the statements that use them (Fig. 5). A *typing context* Γ is a mapping from names to types, where the names can be variables, constants and method names. A *type lookup* is denoted by $\Gamma(x)$, which returns the type of the name x . By $\Gamma[x \mapsto T]$ we denote the update of Γ such that the type of x is set to T . Then $\Gamma[x \mapsto T](x) = T$ and $\Gamma[x \mapsto T](y) = \Gamma(y)$ if $x \neq y$. An over-lined \bar{e} denotes a sequence of syntactic entities e .

The basic idea underlying the typing rules regarding streams in Fig. 5 is that the type `stream<T>` of streams of data items of type T itself *cannot* be defined as a *subtype* of `fut(T)`, since for a future variable x , a query $x?$ gives rise to a Boolean guard whereas for a stream variable x , the query $x?$ is not a Boolean guard because it not only checks whether the stream is empty or not but also whether it has terminated. On the other hand, the type `fut(T)` of futures that refer to return values of type T itself can be defined as a sub-type of `stream<T>` (as specified by the rule T-FUTURESTREAM), where the stream buffer is either empty (denoted by a sentinel \perp) or contains an *infinite* sequence of the particular return value. For such streams

$$\begin{array}{c}
\text{(T-METHOD)} \\
\frac{\Gamma' = \Gamma[\bar{x} \mapsto \bar{T}, \bar{x}' \mapsto \bar{T}'] \quad \Gamma'[\text{destiny} \mapsto \text{fut}(T'')] \vdash s}{\Gamma \vdash T'' \ m(\bar{T} \ \bar{x}) \ \{\bar{T}' \ \bar{x}'; s\}} \\
\\
\text{(T-ASYNCSTREAM)} \quad \text{(T-ASYNCALL)} \quad \text{(T-SYNCALL)} \\
\frac{\Gamma \vdash e : T' \quad \Gamma \vdash \bar{e} : \bar{T} \quad \text{match}(m, \bar{T} \rightarrow T \ \text{stream}, T')}{\Gamma \vdash \mathbf{nd} \ e!m(\bar{e}) : \text{stream}(T)} \quad \frac{\Gamma \vdash e.m(\bar{e}) : T}{\Gamma \vdash e!m(\bar{e}) : \text{fut}(T)} \quad \frac{\Gamma \vdash e : T' \quad \Gamma \vdash \bar{e} : \bar{T} \quad \text{match}(m, \bar{T} \rightarrow T, T')}{\Gamma \vdash e.m(\bar{e}) : T} \\
\\
\text{(T-RETURNSTREAM)} \quad \text{(T-YIELD)} \quad \text{(T-RETURN)} \\
\frac{\Gamma(\text{destiny}) = \text{stream}(T)}{\Gamma \vdash \mathbf{return}} \quad \frac{\Gamma \vdash e : T}{\Gamma \vdash \mathbf{yield} \ e} \quad \frac{\Gamma \vdash e : T}{\Gamma \vdash \mathbf{return} \ e} \\
\\
\text{(T-AWAITSTREAM)} \quad \text{(T-GETSTREAM)} \\
\frac{\Gamma \vdash e : \text{stream}(T) \quad \Gamma \vdash s}{\Gamma \vdash \mathbf{await} \ e? \ \mathbf{finished} \ \{s\}} \quad \frac{\Gamma \vdash e : \text{stream}(T) \quad \Gamma \vdash s \quad \Gamma \vdash x : T}{\Gamma \vdash x = e.\mathbf{get} \ \mathbf{finished} \ \{s\}} \\
\\
\text{(T-FUTURESTREAM)} \\
\frac{\Gamma \vdash e : \text{fut}(T) \quad T \preceq T'}{\Gamma \vdash e : \text{stream}(T')}
\end{array}$$

Fig. 5. Type system.

the **finished** statement never executes, as there is no termination token. Note also that for such infinite streams there is no difference between destructive or non-destructive reads.

We proceed with a brief explanation of the typing rules. A streaming method is well-typed by T-STREAMMETHOD, if its body s is well-typed in the typing context extended by the parameters, local variables, and the return stream. The *destiny* variable in ABS is a local variable which holds a reference to the return stream (or future).

By T-ASYNCALL, an asynchronous method call to a non-streaming method has type $\text{fut}(T)$, if its corresponding synchronous call has type T . Whereas by T-ASYNCSTREAM the type of an asynchronous call of a streaming method is of type $\text{stream}(T)$, if the interface T' of the callee includes the streaming method. As in ABS, by T-SYNCCALL, a call to a method m has type T if its actual parameters have types \bar{T} and the signature $\bar{T} \rightarrow T$ matches a signature for m in the known interface of the callee (given by an auxiliary function *match*). The rule does not allow synchronous call on a streaming method (note the difference between how the function *match* is used in T-ASYNCSTREAM and T-SYNCCALL).

The **yield** statement is well-typed in a streaming method by T-YIELD, if the type of e is T and the enclosing method is a streaming method of type T . **return** statement without parameter is only used in a streaming method to signal the termination of streaming and the method, and is well-typed by T-RETURNSTREAM.

The T-RETURN forces that the expression e of the **return** statement in a non-streaming method is of type T , the return type of the enclosing method.

The **await-finished** is well-typed by T-AWAITSTREAM, if a stream of some type T is awaited and if the statement s is also well-typed. It is not difficult to see how the statement **get-finished** is well-typed by T-GETSTREAM.

3.5. Operational semantics

The operational semantics of the proposed extension is presented below as a transition system in SOS style [11]. First we extend the ABS run-time configuration and then present those rules in the transition system that involve destructive and non-destructive streams.

3.5.1. Runtime configuration

The runtime syntax of ABS is extended by the notion of stream is illustrated in Fig. 6. Configurations cn consist of objects (*object*), invocation messages (*invoc*), futures (*fut*), and data streams (*stream*). The commutative and associative composition operator on configurations is denoted by whitespace. The empty configuration is denoted by ϵ .

The term $ob(o, a, p, q)$ represents an object where o is the object identifier, a assigns values to the object's fields, p is an active process (or idle), and q represents a set of suspended processes.

The term $invoc(o, f, m, \bar{v})$ represents an invocation message, where o is the callee object, f is the identifier of a rendezvous for the return value(s) of the method invocation which can be a stream or a future, depending on the invoked method being a *streaming* method or not, m is the name of the invoked method, and \bar{v} are its arguments.

A process $\{a \mid s\}$ consists of an assignment a of values to the local variables, and a statement s . A process results from the activation of a method invocation in a callee with actual parameters, and an associated future or stream. An **error** is a process where the binding of such method invocation does not succeed.

$cn ::=$	$\epsilon \mid fut \mid stream \mid$	
	$object \mid invoc \mid cn \ cn$	
$object ::=$	$ob(o, a, p, q)$	$p ::=$
	$fut ::= fut(f, value)$	$process \mid \mathbf{idle}$
$stream ::=$	$stream(f, u)$	$a ::=$
		$T \ x \ v \mid a, a$
$process ::=$	$\{a \mid s\} \mid \mathbf{error}$	$value ::=$
	$q ::= \epsilon \mid process \mid q \ q$	$v \mid \perp$
$invoc ::=$	$invoc(o, f, m, \bar{v})$	$u ::=$
		$u.u \mid v \mid \eta \mid \perp$
		$v ::=$
		$o \mid f \mid t \mid (f, n)$

Fig. 6. Runtime configuration.

A future is represented by $fut(f, value)$, where f is the future identifier and $value$ denotes its current value which can either be the actual value returned or \perp which denotes the absence of a return value.

Both destructive and non-destructive streams are semantically represented by $stream(f, u)$, where f is the stream identifier, and u denotes its buffer. Nevertheless, the stream variables referring to destructive streams just hold the stream id, whereas the value of a stream variable referring to a non-destructive stream is a pair (f, n) , where n denotes the associated position in the stream.

The buffer u is a FIFO queue, which contains a sequence of values v , and a special symbol, either \perp which is a sentinel denoting end of buffer, or η which denotes termination of streaming. The \perp is replaced by η after adding the last valid value to the queue when the streaming method terminates. The $u = v.u'$ denotes the head v of the queue u , and its tail u' . In $u' = u.v$, enqueueing the value v to the end of the queue u forms the updated queue u' . The auxiliary function $elem(u, n)$ returns the content at the position n of the sequence u starting from 0.

A value v can be an object identifier, a future or stream identifier, a term t which is a value of a primitive type, or a pair (f, n) which is a value of a variable referring to a non-destructive stream.

Note that all the identifiers in a configuration are unique and terminal: o is used for object, and f both for future and stream identifiers.

The rules of Fig. 8 and 9 operate on the elementary configurations. To have the rules to apply to full configurations, we need the following rule as well:

$$\frac{cn' \rightarrow cn''}{cn \ cn' \rightarrow cn \ cn''}$$

We also use the reduction system proposed in the ABS formal model to evaluate expressions, e.g., $f = \llbracket e \rrbracket_{aol}^{cn}$ in the active process of $ob(o, a, \{l|s\}, q)$ holds if the expression e evaluates to the stream identifier f , in an assignment composed of a and l , where the configuration $cn \ ob(o, a, \{l|s\}, q)$ is given, and cn contains $stream(f, u)$. By definition, $a \circ l(x) = l(x)$ if $x \in dom(l)$ or $a \circ l(x) = a(x)$ otherwise.

The following rule ASYNC-CALL represents asynchronous method invocation in core ABS extended with a check that it is not a streaming method:

$$\frac{\begin{array}{c} (ASYNC-CALL) \\ o' = \llbracket e \rrbracket_{(aol)} \quad \bar{v} = \llbracket \bar{e} \rrbracket_{(aol)} \quad \mathbf{fresh}(f) \quad \neg \mathbf{streamer}(o'.m(\bar{v})) \end{array}}{\begin{array}{c} ob(o, a, \{l|x = e!m(\bar{e}); s\}, q) \\ \rightarrow ob(o, a, \{l|x = f; s\}, q) \quad invoc(o', f, m, \bar{v}) \quad fut(f, \perp) \end{array}}$$

where it sends an invocation message to object o' with the method name m , the future f and the actual parameters \bar{v} . The return value of f is undefined (i.e., \perp). Note that, based on Fig. 6, the definition of v also includes the values f and (f, n) for destructive and non-destructive streams in the extended semantics. Therefore streams can be passed as actual parameters and assigned to formal parameters.

In the rest of this section, we present the semantic rules of the extended ABS, where a data stream is involved. Given in Fig. 7, the rules for the callee side, which only write to the stream, are independent from how the stream is read (i.e., destructively or non-destructively). In the rule YIELD, the active process, which is a streaming method, enqueues the value v to the buffer of the stream f , followed by the sentinel \perp . The rule RETURNSTREAM enqueues the value η to the buffer of the stream f , which is a token denoting termination of streaming values in the buffer.

In the following, the rules for destructive and non-destructive access to the data stream are given. Note that the D and ND are prefixed to the rule names, which stand for the destructive and non-destructive streams, respectively.

3.5.2. Semantics of destructive streams

In the rule D-ASYNCALL, the object o calls asynchronously a streaming method m with arguments \bar{v} on object o' . The return stream is destructive with the fresh identifier f , as the access mode to the return stream of a streaming method is destructive by default. We also use two auxiliary functions in this rule as follows: the function $streamer(o.m(\bar{v}))$ checks if the method $m(\bar{v})$ of the object o is a streaming method. The function $fresh(f)$ guarantees that the newly introduced name f is not already used in the system.

$$\begin{array}{c}
\text{(YIELD)} \\
\frac{v = \llbracket e \rrbracket_{\text{ac}}^{\text{cn}} \quad l(\text{destiny}) = f}{ob(o, a, \{\mid \mathbf{yield} \ e; s\}, q) \quad \text{stream}(f, u, \perp) \rightarrow ob(o, a, \{\mid s\}, q) \quad \text{stream}(f, u, v, \perp)} \\
\text{(RETURNSTREAM)} \\
\frac{l(\text{destiny}) = f}{ob(o, a, \{\mid \mathbf{return} \ s; s\}, q) \quad \text{stream}(f, u, \perp) \rightarrow ob(o, a, \mathbf{idle}, q) \quad \text{stream}(f, u, \eta)}
\end{array}$$

Fig. 7. Operational semantics of streams on the callee side.

$$\begin{array}{c}
\text{(D-ASYNC CALL)} \\
\frac{o' = \llbracket e \rrbracket_{\text{ac}}^{\text{cn}} \quad \bar{v} = \llbracket \bar{e} \rrbracket_{\text{ac}}^{\text{cn}} \quad \text{fresh}(f) \quad \text{streamer}(o'.m(\bar{v}))}{ob(o, a, \{\mid x = e \mid m(\bar{e}); s\}, q) \rightarrow ob(o, a, \{\mid x = f; s\}, q) \quad \text{invoc}(o', f, m, \bar{v}) \quad \text{stream}(f, \perp)} \\
\text{(D-AWAITTRUE)} \\
\frac{f = \llbracket e \rrbracket_{\text{ac}}^{\text{cn}}}{ob(o, a, \{\mid \mathbf{await} \ e? \ \mathbf{finished} \ \{s_1\}; s_2\}, q) \quad \text{stream}(f, v, u) \rightarrow ob(o, a, \{\mid s_2\}, q) \quad \text{stream}(f, v, u)} \\
\text{(D-AWAITFALSE)} \\
\frac{f = \llbracket e \rrbracket_{\text{ac}}^{\text{cn}}}{ob(o, a, \{\mid \mathbf{await} \ e? \ \mathbf{finished} \ \{s_1\}; s_2\}, q) \quad \text{stream}(f, \perp) \rightarrow ob(o, a, \{\mid \mathbf{suspend}; \mathbf{await} \ e? \ \mathbf{finished} \ \{s_1\}; s_2\}, q) \quad \text{stream}(f, \perp)} \\
\text{(D-AWAITTERMINATE)} \\
\frac{f = \llbracket e \rrbracket_{\text{ac}}^{\text{cn}}}{ob(o, a, \{\mid \mathbf{await} \ e? \ \mathbf{finished} \ \{s_1\}; s_2\}, q) \quad \text{stream}(f, \eta) \rightarrow ob(o, a, \{\mid s_1; s_2\}, q) \quad \text{stream}(f, \eta)} \\
\text{(D-GETTRUE)} \\
\frac{f = \llbracket e \rrbracket_{\text{ac}}^{\text{cn}}}{ob(o, a, \{\mid x = e.\mathbf{get} \ \mathbf{finished} \ \{s_1\}; s_2\}, q) \quad \text{stream}(f, v, u) \rightarrow ob(o, a, \{\mid x = v; s_2\}, q) \quad \text{stream}(f, u)} \\
\text{(D-GETTERMINATE)} \\
\frac{f = \llbracket e \rrbracket_{\text{ac}}^{\text{cn}}}{ob(o, a, \{\mid x = e.\mathbf{get} \ \mathbf{finished} \ \{s_1\}; s_2\}, q) \quad \text{stream}(f, \eta) \rightarrow ob(o, a, \{\mid s_1; s_2\}, q) \quad \text{stream}(f, \eta)}
\end{array}$$

Fig. 8. Operational semantics of destructive streams.

The `await` statement in rule D-AWAITTRUE is skipped as there exists a data value v in the buffer. By rule D-AWAITFALSE, the process querying the empty (but not-yet-terminated) stream f will be suspended. To this aim, the statement `suspend` for unconditional suspension is added to the beginning of the sequence of the statements of the process. According to the standard ABS, the `suspend` then suspends the active process, namely, it adds the process to q , where the active object is *idle* and ready to activate a suspended process from q . In rule D-AWAITTERMINATE, the *finished* block s_1 of the statement is selected for execution, since the streaming is terminated, i.e., the head of the buffer of the stream f is equal to the terminating token η .

The rule D-GETTRUE assigns the value v from the head of the stream buffer to the variable x destructively, i.e., v is removed from the buffer. By D-GETTERMINATE, the *finished* block s_1 of the statement is executed followed by s_2 , as the terminating token is observed at the head of the buffer. Note that the state of x remains the same. There is no rule for the `get-finished` statement when the buffer is empty which implies that the active process (and the object) is *blocked* until the buffer contains an element.

3.5.3. Semantics of non-destructive streams

The operational semantics of ABS for those rules that involve non-destructive future-based streams is given in Fig. 9. In ND-ASYNC CALL, an asynchronous call to a streaming method m in o' is given with the actual parameters \bar{v} , that results in a reference to a non-destructive stream. The keyword `nd` denotes the non-destructive access to the resulting stream. Therefore, the return reference to the newly created stream with identifier f is a pair of f and a cursor which is initialized to 0, denoting the first position in the buffer which is initially \perp .

The `await` statement in rule ND-AWAITTRUE is skipped because there is a value v (which is not η) in the buffer of the stream f at the position determined by the cursor of x . By rule ND-AWAITFALSE, the process querying the stream f will

$$\begin{array}{c}
\text{(ND-ASYNCALL)} \\
\hline
\frac{o' = \llbracket e \rrbracket_{\text{acol}} \quad \bar{v} = \llbracket \bar{e} \rrbracket_{\text{acol}} \quad \text{fresh}(f) \quad \text{streamer}(o'.m(\bar{v}))}{ob(o, a, \{\llbracket x = \mathbf{nd} \ e!m(\bar{e}); s \rrbracket, q\}) \rightarrow ob(o, a, \{\llbracket x = (f, 0); s \rrbracket, q\}) \text{in} \text{invoc}(o', f, m, \bar{v}) \text{ stream}(f, \perp)} \\
\hline
\text{(ND-AWAITTRUE)} \\
\frac{(f, n) = \llbracket x \rrbracket_{\text{acol}}^{\text{cn}} \quad v = \text{elem}(u, n)}{ob(o, a, \{\llbracket \mathbf{await} \ x? \ \mathbf{finished} \ \{s_1\}; s_2 \rrbracket, q\}) \text{ stream}(f, u) \rightarrow ob(o, a, \{\llbracket s_2 \rrbracket, q\}) \text{ stream}(f, u)} \\
\hline
\text{(ND-AWAITFALSE)} \\
\frac{(f, n) = \llbracket x \rrbracket_{\text{acol}}^{\text{cn}} \quad \perp = \text{elem}(u, n)}{ob(o, a, \{\llbracket \mathbf{await} \ x? \ \mathbf{finished} \ \{s_1\}; s_2 \rrbracket, q\}) \text{ stream}(f, u) \rightarrow ob(o, a, \{\llbracket \mathbf{suspend}; \mathbf{await} \ x? \ \mathbf{finished} \ \{s_1\}; s_2 \rrbracket, q\}) \text{ stream}(f, u)} \\
\hline
\text{(ND-AWAITTERMINATE)} \\
\frac{(f, n) = \llbracket x \rrbracket_{\text{acol}}^{\text{cn}} \quad \eta = \text{elem}(u, n)}{ob(o, a, \{\llbracket \mathbf{await} \ x? \ \mathbf{finished} \ \{s_1\}; s_2 \rrbracket, q\}) \text{ stream}(f, u) \rightarrow ob(o, a, \{\llbracket s_1; s_2 \rrbracket, q\}) \text{ stream}(f, u)} \\
\hline
\text{(ND-GETTRUE)} \\
\frac{(f, n) = \llbracket y \rrbracket_{\text{acol}}^{\text{cn}} \quad v = \text{elem}(u, n)}{ob(o, a, \{\llbracket x = y.\mathbf{get} \ \mathbf{finished} \ \{s_1\}; s_2 \rrbracket, q\}) \text{ stream}(f, u) \rightarrow ob(o, a, \{\llbracket x = v; y = (f, n + 1); s_2 \rrbracket, q\}) \text{ stream}(f, u)} \\
\hline
\text{(ND-GETTERMINATE)} \\
\frac{(f, n) = \llbracket y \rrbracket_{\text{acol}}^{\text{cn}} \quad \eta = \text{elem}(u, n)}{ob(o, a, \{\llbracket x = y.\mathbf{get} \ \mathbf{finished} \ \{s_1\}; s_2 \rrbracket, q\}) \text{ stream}(f, u) \rightarrow ob(o, a, \{\llbracket s_1; s_2 \rrbracket, q\}) \text{ stream}(f, u)}
\end{array}$$

Fig. 9. Operational semantics of non-destructive streams.

be suspended since the cursor of f denotes the empty position in the buffer (denoted by \perp). By the semantics, it is not difficult to see that this position will contain either a value v or the termination token η . By the rule ND-AWAITTERMINATE, the **finished** block s_1 is selected for execution, as the cursor of f points at a position which contains η .

The rule ND-GETTRUE assigns to the variable x the value v (which is not η) in the stream buffer from the position determined by the cursor of y , and increments the cursor. By ND-GETTERMINATE, the **finished** block s_1 of the statement is selected for execution followed by s_2 , as the cursor of variable y points at the terminating token η in the buffer. Note that the state of x and the cursor are not modified. There is no rule for the **get-finished** statement when the cursor denotes the empty position in the buffer (i.e., \perp) which implies that the active process (and the object) is *blocked*.

3.5.4. Semantics of futures as streams

The type of the value of any expression in ABS at runtime is a subtype of the static type of the expression. The **await** and **get** without **finished** clause can only be applied to futures and Boolean guards, and does exclude the streams. This is guaranteed because **Stream**< T > is not a subtype of **Fut**< T > (discussed in section 3.4). Recall that an **await without finished** clause on a stream is only a syntactic sugar for the one *with* the clause where the following block is empty. Different operational semantics of destructiveness and non-destructiveness does not affect the type system.

In order to support the subtyping relation between $\text{stream}(T)$ and $\text{fut}(T)$ in the operational semantics, as reflected in the type system, we need an extra set of semantic rules, where a future variable appears as the parameter of **await-finished** and **get-finished** statements. This set is presented in Fig. 10. In these rules, only the cases are specified where the future contains a value v or not (empty stream). A resolved future is treated as an infinite stream of the same value v . Therefore, termination of future is not defined. The rule names are prefixed with \mathbb{F} to denote that future appears as a stream.

We can prove on the basis of the operational semantics in a standard manner that all program executions are type-safe, and in our case this additionally ensures proper use of the data streams. This additionally amounts to ensuring that the **await-finished** and **get-finished** constructs are applied at runtime only to the data streams (and futures) and that the **yield** operation is only applied to the context of a streaming method.

3.6. Discussion on buffer size and garbage collection

The buffer of streams can grow indefinitely according to the above semantics. For practical purposes, however, we must take into consideration the finite nature of computer memory. This can be addressed by a different definition of the buffer which is bounded to a maximum size m . The semantics of a successful write operation to a bounded buffer requires a new premise where the buffer size is strictly less than m . If the buffer is full, on the other hand, different design decisions can

$$\begin{array}{c}
\text{(F-AWAITTRUE)} \\
\frac{f = \llbracket e \rrbracket_{\text{act}}^{\text{cn}}}{\text{ob}(o, a, \{\!| \text{await } e? \text{ finished } \{s_1\}; s_2 \}, q) \text{ fut}(f, v) \rightarrow \text{ob}(o, a, \{\!| s_2 \}, q) \text{ fut}(f, v)} \\
\text{(F-AWAITFALSE)} \\
\frac{f = \llbracket e \rrbracket_{\text{act}}^{\text{cn}}}{\text{ob}(o, a, \{\!| \text{await } e? \text{ finished } \{s_1\}; s_2 \}, q) \text{ fut}(f, \perp) \rightarrow \text{ob}(o, a, \{\!| \text{suspend; await } e? \text{ finished } \{s_1\}; s_2 \}, q) \text{ fut}(f, \perp)} \\
\text{(F-GETTRUE)} \\
\frac{f = \llbracket e \rrbracket_{\text{act}}^{\text{cn}}}{\text{ob}(o, a, \{\!| x = e. \text{get } \text{ finished } \{s_1\}; s_2 \}, q) \text{ fut}(f, v) \rightarrow \text{ob}(o, a, \{\!| x = v; s_2 \}, q) \text{ fut}(f, v)}
\end{array}$$

Fig. 10. Semantics of futures as streams.

be made. For instance writing to a full buffer can be blocking, i.e., the process is blocked until the buffer size is less than m , or it is non-blocking but signals the process about the failure. A successful destructive read operation, on the other side, decrements the buffer size allowing the buffer to shrink. However this is not the case for non-destructive streams, as the non-destructive read cannot change the buffer size since the data will possibly be read by other cursors. Hence, we need a garbage collection mechanism (GC) for non-destructive streams.

By definition, destructive streams do not cause any garbage. However, we can have a definition of garbage for non-destructive streams. In this section, garbage means a data element in the buffer of a non-destructive stream, which is read by *all* the *existing* cursors. In what follows, we define a GC that is executed periodically. It first obtains all the existing cursors in the system and then collects the garbages accordingly. Some of the cursors can be obtained from the immediate value of a variable, while other cursors can be wrapped with an outer future or stream in a nested way. For instance, if the future variable $x : \text{fut}(\text{stream}(T))$ is resolved can possibly contain a cursor. To include these cursors in the GC, first we need some definitions. Let type T denote either a primitive type P (a type that is not a future nor a stream) or a non-primitive type N as follows:

$$T ::= P \mid N$$

$$N ::= \text{stream}(T)$$

For notational convenience we rewrite the type $\text{fut}(T)$ to $\text{stream}(T)$. We also rewrite a run-time future object to a destructive stream with at most one element so that it can be typed as a stream. Therefore a future $\text{fut}(f, \perp)$ is rewritten to $\text{stream}(f, \perp)$ and $\text{fut}(f, v)$ to $\text{stream}(f, v\eta)$ in cn . The following algorithm obtains the set of all cursors in the system, based on which it marks the garbage:

1. for each object $(o, a, \{\!| o | s_0 \}, \{\!| l_1 | s_1 \}, \dots, \{\!| l_k | s_k \})$ in the system, the set of cursors that can be obtained in object o : $\text{cursor}_o = \{\text{cursor}_o(\llbracket x \rrbracket_{\text{act}}^{\text{cn}} : T) \mid x \in a \cup \bigcup_{i=0}^k l_i\}$ where $\text{cursor}_o(v : T)$ returns all the existing cursors obtained from value v of type T in object o . The set of all the cursors existing in the system: $\text{cursors} = \bigcup \text{cursor}_o$. Note that for simplicity we assume there is no name conflict of variable names in the mappings.
2. for each stream identity f in the system, $\text{min}_f = \min(\{n \mid (f, n) \in \text{cursors}\})$, where $\text{min}(S)$ returns the smallest number in a set S of numbers.
3. for each $\text{stream}(f, u)$ in the system, all the data elements in u with the index less than min_f are garbage and must be collected.

For simplicity we use $v : T$ for typing value v instead of using the formal run-time type system. Below we define $\text{cursor}_o^n(v : T)$ inductively, where n is the number of times the term stream appears in the type T of the value v , e.g., n is 0, 1, and 2 for the type P , $\text{stream}(P)$ and $\text{stream}(\text{stream}(P))$, respectively. Note that n is finite, as the type of a variable is a string with a finite length.

Base case: a cursor can neither be obtained from a value with a primitive type (step 0), nor a value that refers to a destructive stream or a future of a primitive type P (step 1). Whereas one cursor can be obtained from a value that refers to a non-destructive stream of a primitive type P (step 1).

$$\begin{array}{ll}
\text{cursor}_o^0(v : P) = \emptyset & \\
\text{cursor}_o^1(f : \text{stream}(P)) = \emptyset & \text{where } \text{stream}(f, u) \in \text{cn} \\
\text{cursor}_o^1((f, n) : \text{stream}(P)) = \{(f, n)\} & \text{where } \text{stream}(f, u) \in \text{cn}
\end{array}$$

$$\begin{array}{c}
\text{(T-STREAM)} \\
\frac{\Delta(f) = \text{stream}\langle T \rangle \quad \forall i \in [1..n]. (val_i \notin \{\perp, \eta\} \Rightarrow \Delta(val_i) = T)}{\Delta \vdash_R \text{stream}(f, (val_1, \dots, val_n)) \mathbf{ok}}
\end{array}
\qquad
\begin{array}{c}
\text{(T-STATESTREAM)} \\
\frac{\Delta(val) = (\text{stream}\langle T \rangle, \text{Nat}) \quad \Delta \vdash_R v : \text{stream}\langle T \rangle}{\Delta \vdash_R \text{stream}\langle T \rangle \ v \ \text{val} \ \mathbf{ok}}
\end{array}$$

$$\begin{array}{c}
\text{(T-INVOCSTREAM)} \\
\frac{\Delta(f) = \text{stream}\langle T \rangle \quad \Delta(\bar{v}) = \bar{T} \quad \text{match}(m, \bar{T} \rightarrow \text{stream}\langle T \rangle, \Delta(o))}{\Delta \vdash_R \text{invoc}(o, f, m, \bar{v})}
\end{array}$$

Fig. 11. The typing rules of streams for run-time configurations.

Inductive step: the induction hypothesis is that $\text{cursor}_0^n(v : N)$ returns all the cursors obtained from the value v of type N where $n = k$. Below we show how we obtain the cursors obtained from a value for $n = k + 1$ using the hypothesis:

$$\begin{aligned}
\text{cursor}_0^{k+1}(f : \text{stream}\langle N \rangle) &= \bigcup_{v_i \in s_n(u)} \text{cursor}_0^k(v_i : N) \\
&\quad \text{where } \text{stream}(f, u) \in cn \\
\text{cursor}_0^{k+1}((f, n) : \text{stream}\langle N \rangle) &= \bigcup_{v_i \in s_n(u)} \text{cursor}_0^k(v_i : N) \cup \{(f, n)\} \\
&\quad \text{where } \text{stream}(f, u) \in cn
\end{aligned}$$

where $s_n(u)$ denotes a set of elements in the buffer u of $\text{stream}(f, u)$ with index greater or equal to n , except special elements \perp and η .

In order for the above algorithm to work, every data element in the buffer u must have an absolute index starting from zero for the first element added to the buffer. Recall that, by the semantics of ABS, every synchronous and asynchronous method call forms a (suspended or active) process in the called object which is denoted by $\{i|s_i\}$. Hence, the step 1 covers all the variable assignments in an object (cursor_0) and subsequently in the whole system (cursors). Also note that there is no need to distinguish the cursors by their variable names or their processes or objects. The only relevant aspect of the cursor for GC is that there exists at least one cursor that points at a specific index of the buffer.

4. Subject reduction for the extended ABS

A *run* is a sequence of transitions from an *initial state* based on the rules of the operational semantics, where *initial state* consists of $ob(\text{start}, \epsilon, p, \emptyset)$, an initial object, start, with only one process p that corresponds to the main block of the program. The subject reduction for ABS is already proven in [3], namely, it is shown that a run from a well-typed initial configuration will maintain well-typed configurations, particularly, the assignments preserve well-typedness and method bindings do not give rise to the **error** process. In this section, we aim to extend the proof for the ABS subject reduction theorem to also include the notion of stream as specified in this paper.

The typing context for the run-time configurations Δ extends the static typing context Γ with typing dynamically created values (entities created at run-time), namely, object and future identifiers. Let $\Delta \vdash_R cn \ \mathbf{ok}$ express that the configuration cn is well-typed in the typing context Δ . The typing rules for run-time configurations are defined for ABS and extensively discussed in [3]. The newly added rules for typing streams are shown in Fig. 11. By T-STREAM, the stream f is of type $\text{stream}\langle T \rangle$ if the buffer only contains values of type T or the special tokens η and \perp . By T-STATESTREAM, a variable v that refers to a stream val and provide non-destructive access to it is well-typed. Nat denotes the type of natural numbers. The type of val is a pair of the stream type and a Nat that holds the cursor to the stream. The rule T-INVOCSTREAM allows the return type of an asynchronous method invocation to be a stream as well.

In [3] (1) it is proven that the initial object corresponding to the main block of a well-typed program is well-typed and also (2) it is shown that the well-typedness of runtime configuration is preserved by reductions (Theorem 1). The proof for (1) also applies here. We only need to extend the proof for (2) with respect to the new transition rules introduced in section 3 as follows.

Theorem 1 (Subject Reduction). *If $\Delta \vdash_R cn \ \mathbf{ok}$ and $cn \rightarrow cn'$, then there is a Δ' such that $\Delta \subseteq \Delta'$ and $\Delta' \vdash_R cn' \ \mathbf{ok}$.*

Proof. The proof is by induction over the defined transition rules in the operational semantics. We assume objects, futures, streams and messages not affected by a transition remain well-typed, and are ignored below. The auxiliary function $\text{match}(m, \bar{T} \rightarrow \text{stream}\langle T \rangle, T')$ checks if a method m with $\bar{T} \rightarrow \text{stream}\langle T \rangle$ is provided by the interface T' .

- **Process Suspension.** It is immediate that the rules D-AWAITTRUE, ND-AWAITTRUE, F-AWAITTRUE, D-AWAITFALSE, ND-AWAITFALSE, F-AWAITFALSE D-AWAITTERMINATE, ND-AWAITTERMINATE, D-GETTERMINATE and ND-GETTERMINATE preserve the well-typedness.
- **YIELD.** By assumption, we have $\Delta \vdash_R ob(o, a, \{\!| \mathbf{yield} \ e; s; \!\}, q) \ \mathbf{ok}$, $\llbracket e \rrbracket_{aol} = v$ and $\Delta \vdash_R stream(f, u.\perp) \ \mathbf{ok}$. Obviously, $\Delta \vdash_R ob(o, a, \{\!| s; \!\}, q) \ \mathbf{ok}$. Since $l(\text{destiny}) = f$ and l is well-typed, we know that $\Delta(\text{destiny}) = \Delta(f)$. Let $\Delta(f) = stream(T)$. By T-YIELD, $\Delta \vdash_R e : T$ and subsequently $\Delta(v) = T$, so $\Delta \vdash_R stream(f, u.v.\perp) \ \mathbf{ok}$.
- **RETURNSTREAM.** By assumption, we have $\Delta \vdash_R ob(o, a, \{\!| \mathbf{return}; s; \!\}, q) \ \mathbf{ok}$, and $\Delta \vdash_R stream(f, u.\perp) \ \mathbf{ok}$. Obviously, $\Delta \vdash_R ob(o, a, \{\!| s; \!\}, q) \ \mathbf{ok}$ and $\Delta \vdash_R stream(f, u.\eta) \ \mathbf{ok}$.
- **D-ASYNCALL.** Let $\Delta \vdash_R ob(o, a, \{\!| x = e!m(\bar{e}); s; \!\}, q) \ \mathbf{ok}$. We first consider the case $e \neq \mathbf{this}$. By T-ASYNCSTREAM, we may assume that $\Delta \vdash e!m(\bar{e}) : stream(T)$ and by T-ASSIGN that $\Delta(x) = stream(T)$. Therefore, $\Delta \vdash e : T'$ and $\Delta \vdash \bar{e} : \bar{T}$ such that $match(m, \bar{T} \rightarrow T \ stream, T')$. Assume that $\llbracket e \rrbracket_{aol} = o'$ and let $\Delta(o') = C$ for some class C . Based on [3], there is a Δ' such that $\Delta' \vdash_R \llbracket e \rrbracket_{aol} : T'$ and $\Delta'(o') = C$, so $C \leq T'$. By assumption class definitions are well-typed, so for any class C that implements interface T' we have $match(m, \bar{T} \rightarrow T \ stream, C)$. Also $\llbracket \bar{e} \rrbracket_{aol}$ similarly preserves the type of \bar{e} . Let $\Delta'' = \Delta'[f \mapsto stream(T)]$. Since $fresh(f)$ we know that $f \notin dom(\Delta')$, so if $\Delta' \vdash_R cn \ \mathbf{ok}$, then $\Delta'' \vdash_R cn \ \mathbf{ok}$. Since $\Delta' \vdash e!m(\bar{e}) = \Delta''(f)$, we get $\Delta'' \vdash_R ob(o, a, \{\!| x = f; s; \!\}, q) \ \mathbf{ok}$. Furthermore, $\Delta'' \vdash invoc(o', f, m, \bar{v}) \ \mathbf{ok}$ and $\Delta'' \vdash_R stream(f, \perp) \ \mathbf{ok}$. The case $e = \mathbf{this}$ is similar, but uses the class of \mathbf{this} directly for the match (so internal methods are also visible).
- **ND-ASYNCALL.** Let $\Delta \vdash_R ob(o, a, \{\!| \mathbf{nd} \ x = e!m(\bar{e}); s; \!\}, q) \ \mathbf{ok}$. The argument is similar to the above case, but we get $\Delta'' \vdash_R ob(o, a, \{\!| x = (f, 0); s; \!\}, q) \ \mathbf{ok}$ as the consequence, in addition to $\Delta'' \vdash invoc(o', f, m, \bar{v}) \ \mathbf{ok}$ and $\Delta'' \vdash_R stream(f, \perp) \ \mathbf{ok}$.
- **D-GETTRUE.** By assumption, $\Delta \vdash_R ob(o, a, \{\!| x = e.\mathbf{get \ finished}\{s_1; s_2\}; \!\}, q) \ \mathbf{ok}$, $\Delta \vdash_R stream(f, v.u) \ \mathbf{ok}$, and $\llbracket e \rrbracket_{aol} = f$. Let $\Delta(f) = stream(T)$. Consequently, $\Delta \vdash_R e.\mathbf{get \ finished}\{s_1; s_2\} : T$ and $\Delta(v) = T$, so $\Delta \vdash x = v$, $\Delta \vdash_R ob(o, a, \{\!| x = v; s; \!\}, q) \ \mathbf{ok}$ and $\Delta \vdash_R stream(f, u) \ \mathbf{ok}$. A similar argument applies for F-GETTRUE where f is the identity of a future object $fut(f, v)$.
- **ND-GETTRUE.** By assumption, $\Delta \vdash_R ob(o, a, \{\!| x = y.\mathbf{get \ finished}\{s_1; s_2\}; \!\}, q) \ \mathbf{ok}$, $\Delta \vdash_R stream(f, u) \ \mathbf{ok}$, $\llbracket y \rrbracket_{aol} = (f, n)$ and $elem(u, n) = v$. Let $\Delta(f) = stream(T)$. Consequently, $\Delta \vdash_R y.\mathbf{get \ finished}\{s_1; s_2\} : T$ and $\Delta(v) = T$, so $\Delta \vdash x = v$, $\Delta \vdash_R ob(o, a, \{\!| x = v; s; \!\}, q) \ \mathbf{ok}$, and $\Delta \vdash y = (f, n + 1) \ \mathbf{ok}$. \square

5. Data streams in distributed systems

In [12] a scalable distributed implementation of the ABS language is described. In this section we adapt our proposed notion of data streams in ABS to reduce the possible overhead of data steaming in a distributed setting.

To this aim, each streaming method is enabled to package the return values, that is, the method populates its return stream buffer possibly not once per value, but once per sequence of values. The package size can be specified explicitly as a parameter or can be selected based on the underlying deployment, e.g., it can be equal to the packet size of the TCP/IP technology involved. As such the number of packets to be transferred through the network is minimized.

There are two conditions when the package is streamed *before* its size is equal to the pre-specified package size: 1) when the streaming method terminates; 2) when the streaming method cooperatively releases control. The first condition is obvious, while the second prevents a specific kind of deadlock configuration. In general, ABS programs may give rise to deadlocks (see [8] for a discussion of deadlock analysis of ABS programs). However the notion packaging data streams should not give rise to additional deadlock possibilities.

The above second condition prevents the following kind of deadlock situation. Note that *package size* = n means that the number of yielded values needs to be equal to n , so that they are streamed as a package, except for the last package where the size may be less than n . Suppose there are two objects o_1 and o_2 in the run-time configuration where o_1 executes an active process which corresponds to method m_1 given by

$$m_1() \{ r = o_2!m_2(); \mathbf{await} \ r?; o_2!\mathbf{satisfier}() \}$$

and the specification of the streaming method m_2 is an active process p in object o_2 given by

$$m_2() \{ \mathbf{yield} \ x; \mathbf{await} \ e; \mathbf{yield} \ y; \}$$

Furthermore, suppose the method *satisfier* in o_2 changes the object state so that the expression e (which is *False* initially) evaluates to *True*. It is not difficult to see that for all $n \geq 2$, where n is the package size of the stream, the run-time configuration is deadlocked. The reason is that the first yielded value is not streamed before p releases control, as the package size is smaller than n . The deadlock possibility can be generalized to a category of programs where a streaming method releases control before it communicates the values which are yielded. The solution is that the package with the size smaller than n is streamed, before the process cooperatively releases control or blocks.

6. Implementation

In this section, we present a prototype implementation of future-based data streams as an API written in ABS. This API (see Fig. 12) can be used to simulate the semantics of data streaming in ABS itself. The implementation details of the API can be found online.¹

As discussed in section 3, the `Stream<T>` datatype is parametrically polymorphic in its contained values of type T . The original ABS specification, however, offers besides parametric polymorphism also *subtype* polymorphism, through its interface types. In general, when defining and implementing languages with support for subtype-polymorphism, often the issue of *variance* arises: where in the code it is allowed (i.e. type-safe) to upcast to a supertype or downcast to a subtype. For example, given a subtype relation (T is subtype of U), a structure S is called covariant if $S<T>$ is safe to “upcast” to $S<U>$; contravariant if safe to “downcast” $S<U>$ to $S<T>$; invariant if none of the above two hold, i.e. subtype polymorphism cannot be used for this structure, but other methods of polymorphism (e.g. parametric) perhaps can. In practice, the “rule of thumb” suggests that structures which are exclusively read-only (i.e. immutable) are allowed to be covariant, structures that are written-only (e.g. log files) contravariant, and structures that are read-write must be invariant.

The extension of ABS with stream that we describe in this paper, strictly separates at the syntax level the role of the producer of values (write to the stream structure) with the role of the consumer (read from the stream). Since the producer can only append (produce) new values to the stream and not alter (mutate) past values, from the sole point of the consumer the stream structure seems as “immutable” (covariant). In this sense, a consumer holding a variable of type `Stream<T>` should be allowed to upcast it to type `Stream<U>`. Conversely, the producer is allowed to `yield` values of subtype T , if the method call’s return type is typed as `Stream<U>`. As such, at the surface level (syntax and type system) it is acceptable for the `Stream` structure to be treated as covariant; however, at the implementation level it still remains a challenge on how to guarantee type safety at the host language (in our case, Haskell).

The Haskell language has parametric polymorphism but lacks built-in support for subtype polymorphism; for this reason, the ABS-Haskell backend compiler generates dynamic “upcasting” function calls where needed. However, this technique cannot be applied as well with Haskell’s builtin `vector` datatype, which is a low-level built-in structure that cannot be made covariant or contravariant since it has been fixed-byte allocated in memory heap upon creation. For this reason, and also the fact that arrays are in general a mutable (read-write) data structure, the vectors in ABS (borrowed from Haskell) are treated as invariant. Since the implementation of streams in ABS relies currently on vectors, there is the practical limitation of having the `Buffer` type to be invariant. Similarly the `Stream` and `Fut` datatypes are treated as invariant, because the ABS-Haskell backend treats each future `Fut<T>` as a pointer to a vector size-1 stored in the heap that holds the value of T . Based on this practical limitation, the `Stream<T>` datatype introduced in this extension to ABS is treated as subtype-invariant, with support for parametric polymorphism.

The API is semantically compliant with the semantics of data streams defined in this paper: The method that yields to a stream is separated from the access mode of readers to the stream (i.e., either destructive or non-destructive). Every reader has access to a stream via an instance of either `Dref` or `NDref` for destructive or non-destructive access mode, respectively. Furthermore a stream variable (that refers to an instance of `Dref` or `NDref`) is only typed by the `Stream` interface, abstracting from the underlying access mode.

The interface `Buffer<T>` is implemented by the class `CBuffer`. The FIFO buffer (an instance of `CBuffer`) is implemented by a vector whose elements are of type `Maybe<T>`, namely, each of which contains either a value (`Just(v)` where v is of type T) or `Nothing`. A position in the vector can have three different states: It contains `Just(v)` (a value v that can be read), `Nil` (the position is empty and will be filled), and `Nothing` (a token of type `Maybe<T>`) that denotes termination of the stream. The interface `Buffer<T>` provides the methods `yield()` and `terminate()` to the streaming method in order to write to the buffer and to explicitly terminate the stream of data values, respectively. The termination enqueues `Nothing` to the buffer and is meant to be the last statement in the definition of the streaming method (to simulate the terminating **return**). A stream maintains a global index `wrt` to the buffer which denotes the position where the next yielded value is written. It is incremented by every time calling `yield`. In destructive read, the `CBuffer` maintains a global index (i.e., `rd`) to the buffer for all the readers of the stream, whereas in non-destructive read, every reader (i.e., `NDref` instance) maintains a local index (i.e., `cursor`) to the buffer.

The reader can read from a stream by asynchronously calling `pull()` on the `Stream` object that returns a *future* representing the next data value, whether resolved or not. The operation `pull` is overridden in `Dref` and `NDref` for destructive and non-destructive read from the buffer, respectively. The former calls `dread()` method of the `Buffer` which returns the first valid element in the vector, indicated by the index `rd` in buffer, and increments `rd`. Whereas the latter calls `ndread(cursor)` of the `Buffer` where the `cursor` is a field of the `NDref`, which returns the element indicated by the index `cursor` in the vector. The implementation of `dread()` and `ndread(cursor)` is given in Fig. 13 where **await** at lines (4) and (8) cooperatively release control until the condition (indicating whether the buffer element has been produced) holds.

¹ <https://github.com/kazadbakht/ABS-Stream/blob/master/lib/Streams.abs>.

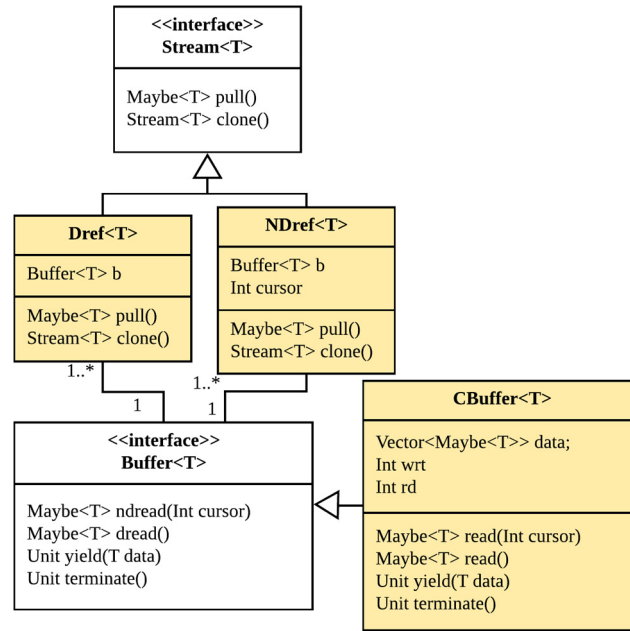


Fig. 12. Class diagram of ABS Stream library.

```

(1) Maybe<T> dread() {
(2)   Int temp = rd;
(3)   rd = rd + 1;
(4)   await (buffer[temp] != Nil);
(5)   return buffer[temp]; // is not null
(6) }

(7) Maybe<T> ndread(cursor) {
(8)   await (buffer[cursor] != Nil);
(9)   return buffer[cursor]; // is not null
(10) }
  
```

Fig. 13. Destructive and non-destructive read in CBuffer.

Also the method `clone` is used to copy a non-destructive stream object, a new instance of `NDref` which has a reference to the same stream but a new cursor which is initialized with the value of the cursor of the original object. For destructive streams, the method only returns the reference to **this** which is of type `Dref`.

Awaiting the future resulting from calling `pull()` queries the availability of next data value. Therefore, statement `await r? finished {S}` is expressed in the library as follows:

```

(1) f = r!pull(); // a future f to the target data value
(2) await f?; // awaits if the future f is not resolved yet
(3) m = f.get; // gets the resolved data value
(4) if (m == Nothing) // Nothing is the special token denoting the termination
(5) {S}
  
```

where `r` is a reference to a stream object and `S` is a block of statements. This can either give rise to the release of control in case the data is not available (line 2) or to skip otherwise. The variable `m` is of type `Maybe<T>` which contains either the value `v`, denoted by `Just(v)`, where `v` is of type `T`, or `Nothing`.

Similarly, statement `x = r.get finished {S}` can be expressed using the library as follows:

```

(1) f = r!pull(); // a future to the target data value
(2) m = f.get; // gets the resolved data value
(3) case (m) {
(4)   Nothing => {S} // "Nothing" is the special token denoting the termination
(5)   Just(v) => {x = v} // the value v is assigned to x
(6) }
  
```


In line 2, the object running this process blocks until the data value is written to the future f .

The keyword `nd` is implemented in the API by a Boolean argument passed to the called streaming method. The argument specifies whether the return object of the streaming method to be an object of class `Dref` or class `NDref`.

The following snippet shows how the library is used to stream *integer* data values. The streaming method `m` instantiates a stream, delegates yielding values to the stream asynchronously to an auxiliary method `m2`, and returns the stream to the caller. Note that `m` sends the same list of parameters it receives to `m2`.

```
// caller :
// False means the return stream is non-destructive
Fut<Stream<Int>> f = o!m(False, ...);
Stream<Int> r = f.get; // r is a reference to the stream
// reading from the non-destructive stream r
...

// callee
Stream<Int> m(Bool isDestr, ...){
    Buffer<Int> b = new Cbuffer();
    // b is the return stream which is filled by m2
    this!m2(b, ...);
    // isDestr determines the access mode to the stream
    // i.e., destructively or non-destructively
    if (isDestr)
        return new Dref<Int>(b);
    else
        return new NDref<Int>(b, 0);
}

// the implementation of the callee
Unit m2(Buffer<Int> b, ...) {
    Int x ;
    ...
    b!yield(x); // yield a value
    ...
    b!terminate(); // termination token
}
```

Remark. In the above-mentioned API, having multiple readers for one stream may result in a performance bottleneck, as the buffer object itself queries the availability of data item to be returned for every pull request (via `dread` or `ndread`). Alternatively, such availability check can be delegated to the reader's pull method itself. However in the current design, doing the checks of line (4) and (8) in Fig. 13 in their corresponding pull definitions give rise to busy-wait polling. The key feature that enables delegating the check without busy-wait polling is the data type `Promise<T>`. A promise is of this type is either contains a data value of type `T` (resolved) or not. An unresolved promise p can be resolved by $p.give(v)$ with some data value v . Similar to futures, `get` and `await` operations can also be applied to promises.

By converting the type of the buffer vector from `Maybe<T>` to `Promise<Maybe<T>>`, the methods `dread` and `ndread` can immediately return to the reader's pull request the promise object in the vector that holds the expected value, which is either already resolved and can be retrieved or will be resolved in the future. In this new design, the `ndread` in `CBuffer` only returns the promise without availability check as follows:

```
(1) Promise<Maybe<T>> ndread(cursor) {
(2) return buffer[cursor];
(3) }
```

And pull method in `NDref` which checks the availability is given as follows:

```
(1) Maybe<T> pull() {
(2) Fut<Promise<Maybe<T>>> f = b!ndread(cursor); // cursor is a field in NDref
(3) Promise<Maybe<T>> p = f.get;
(4) await p?; // availability check is moved to pull
(5) Maybe<T> result = p.get;
(6) case (result) {Just(v) => cursor = cursor + 1;}
(7) return result;
(8) }
```

In line (6), if `result` is equal to `Nothing` then `cursor` is not incremented, such that the next pull requests to this object result in `Nothing`. Similar changes apply to pull in `Dref` and `dread()` in `CBuffer` for destructive read, except the index `rd` which is updated in `dread`.




Fig. 14. Execution time of parallel map-reduce program for 10^6 data values.

An implementation similar to the example in Fig. 2 using the above API is provided online.² We also ran the implementation on a PC which was an Intel Core i7-5600U 2.60GHz \times 4 with 12GB RAM, and 64-bit Ubuntu 16.04 LTS as the operating system. In Fig. 14, we represent the time measured for execution of the program for different number of parallel processors in 5 runs. On average for this specific implementation, we observed that for two cores we achieve $1.59\times$ speedup compared to one core, and for four cores $1.13\times$ speedup compared to two cores.

Lower speedup achieved for higher number of processors, among other reasons, stems from the fact that a stream is a shared resource where parallel access and yielding values to it in a safe manner is limited. This is confirmed experimentally, by adding more workload on the reader per data value that it reads from the stream. As such, the access rate to the stream becomes low enough such that the stream is not a performance bottleneck. With this modification, we could achieve up to 1.41 speedup for four cores compared to two cores.

The current implementation does not feature garbage collection of streams: produced data values are stored in a vector which dynamically (at-runtime) grows indefinitely (until memory exhaustion). This choice was made for a separation of concerns. This orthogonal issue of garbage collection can be trivially solved in the case of destructive streams: all produced values before the global index can be considered as garbage. A future implementation should automatically reclaim the space for such values and appropriately resize (shrink) the vector. In the case, however, of non-destructive streams, some extra bookkeeping and communication is involved to have safe, distributed garbage collection of streams. One possible implementation would require storing at the producer's side a global (system-wide) minimum of all the readers' local cursors. Besides this bookkeeping of the producer, once a reader forwards a `Stream<T>` to another ABS process (local or remote), it involves notifying the producer about the local minimum of the new reader process. Furthermore, in case of a real distributed system, the producer should monitor the quality of the network connection to every reader, otherwise it runs the risk of memory leaking from a dropped connection to a reader.

7. Case study

Simulation of massive social networks is of great importance. Typically an issue in this context is that larger networks are structurally different from smaller networks generated based on the same models [13]. Analysis of social networks is relevant to many scientific domains, e.g., data-mining, network sciences, physics, and social sciences [14]. In this section, we investigate social network simulation based on so-called *preferential attachment* (PA) [15].

Modeling and implementation of preferential attachment is for standard ABS already extensively investigated both for multicore [16] and distributed [12] architectures. Here we focus on how the novel notion of destructive data streams influences the performance of the system presented in [12]. To adopt data streams, we have modified the communication pattern of the active objects, where instead of one request per message, a batch of requests is sent to an active object via one invocation of a streaming method and where the return values are streamed to the caller via data streaming. The performance gain, discussed below, can be attributed almost entirely to the batching responses replacing the sending of one packet per return value. Note, such a batching mechanism is integrated naturally in the context of data streams.

7.1. Problem definition

We represent a social network by means of a graph, where the nodes and the edges of the graph denote, respectively, the members of the network and the connections between them. Every node has a *degree* which is the number of edges that is connected to the node (the number of connection with the members of the network). A social network graph based on PA is generated by adding a sequence of nodes preferentially to the existing graph, starting with an initial graph. The preference is the degree of the nodes of the existing graph, that is, in order to add a new node to the existing graph, the

² https://github.com/kazadbakht/ABS-Stream/blob/master/Examples/map_reduce.abs.




Fig. 15. Array representing the graph.

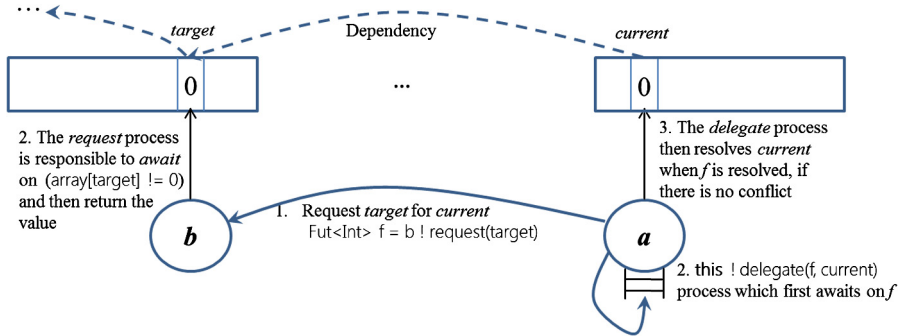


Fig. 16. Dealing with unresolved dependencies.

higher the degree of an existing node, the higher the probability to make a connection with the new node. Adding each new node changes the preference for the next new nodes, as the overall degree of the graph changes.

The problem at hand has two main parameters: n represents the target number of nodes of the graph, and m represents the number of distinct edges that a new node makes with the existing graph. We use q to denote the size (i.e., the number of nodes) of the initial graph. Having the above parameters and the initial graph, the number of edges of the final graph is known in advance. As shown in Fig. 15, the graph is represented by an array, inspired by the so-called copy model of [17]. The slots (i.e., indices) regarding node u , where $q < u \leq n$, represent the m edges connecting the new node u to the existing graph relative to the node u , that is, the graph with the nodes $1, \dots, u - 1$. Each edge is represented by a pair of array slots, denoting the two ends of the edge. The values for the unresolved slots of u in the array, denoted by the question marks, are determined by uniformly selecting m slots with distinct values from all the slots positioned before the slots regarding the node u , namely, the slots regarding the existing graph with the nodes $1, \dots, u - 1$. The uniform selection of the slots respects the preference for each node mentioned before, since the degree of each node is equal to the number of occurrences of the node in the array.

The PA-based graph, with the parameters n , m , and the initial graph given, is generated by resolving the unresolved slots of the array representing the graph. A sequential algorithm is fairly straightforward: It starts from the node $q + 1$ to node n resolving the slots. The parallel and distributed algorithms, however, require partitioning the array and resolving the partitions in parallel, which may give rise to unresolved dependencies, that is, resolving a slot requiring the value of yet another unresolved slot. Any parallel algorithm needs to preserve such dependencies until they are resolved, in order to respect the PA model.

Fig. 16 illustrates the high-level scheme to deal with the unresolved dependencies proposed in [12], using the notion of an active object, future, and cooperative scheduling in ABS. ABS is naturally suited for the PA-based network generation because of distributed nature of the solution. It also improves the algorithm by avoiding low-level synchronization mechanisms thanks to utilizing a powerful programming model and programming constructs. For instance, by leveraging cooperative scheduling, the solution proposed in this section abstracts from maintaining explicit bookkeeping of the requests and their corresponding responses (see below).

In above scheme, the current, unresolved slot which belongs to the active object a requires the value of the unresolved target slot which belongs to the active object b . To this aim, the object a asynchronously calls the *request* method of the object b , and delegates the resulting future as a suspended process in its queue, so that the active process continues with the rest of its partition. On the other side, the request awaits on a Boolean condition which checks if the target is resolved and returns the value. Finally, the *delegate* method which awaited on the future, gets the future value and processes it. Processing the value means resolving the target if there is no conflict. A conflict happens if the new value makes a multi-edge, i.e. more than one edge between two nodes.

7.2. Incorporating data streams

The generation of distributed PA-based graphs as described above is fairly high-level and intuitive at the modeling level. However, the number of messages and return values communicated among the active objects poses a considerable overhead. Packaging the requests and the corresponding return values can considerably improve the performance of the run-time system.

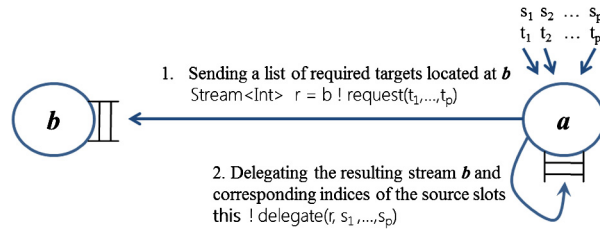


Fig. 17. The modified approach using destructive streaming.

```

1: Each actor  $o$  executes the following in parallel
2: Unit run(...)
3: while the whole partition is not yet processed do
4:   /*
      Resolve the slots. Next pack of unresolved sources =  $[s_1, \dots, s_p]$ 
      from the partition belonging to this object, and its corresponding
      targets =  $[t_1, \dots, t_p]$  whose owning partition hosted by
      some object  $w$  are calculated
5:   */
6:   Stream<Int>  $f = w ! \text{requests}(\text{targets});$ 
7:   this !  $\text{delegates}(f, \text{sources});$ 
8:
9:   Int stream requests(List<Int> targets)
10:  while targets is not nil do
11:    Int tar = head(targets);
12:    targets = tail(targets);
13:    await (arr[tar]  $\neq 0$ );
14:    yield arr[tar];
15:
16:  Unit delegate(Stream<Int> r, Int sources)
17:  while True do
18:    Int val;
19:    await r? finished {break;}
20:    val = r.get;
21:    Int src = head(sources);
22:    sources = tail(sources);
23:    // Use val to resolve arr[src]

```

Fig. 18. The sketch of the data streaming in the modified approach.

In the distributed scheme in Fig. 16, the request is sent per each required target slot, which is too fine-grained. Instead, we propose a modification of the algorithm so that the requests for the target values located on the same active object are sent together as a package of requests via one message, and the returning values are received via a stream with packaging capability.

Dealing with unresolved dependencies in the modified approach is shown in Fig. 17. For all $i \in [1, p]$, each pair (s_i, t_i) represents a request from object a to object b , where s_i represents the index of an unresolved slot belonging to the partition hosted by a , and t_i represents its corresponding slot belonging to the partition hosted by b . The value obtained from each t_i is used to resolve the unresolved slot s_i . Assuming p is the package size, the list $[t_1, \dots, t_p]$ is sent to b as a package of requests, and the requests process returns corresponding values per each t_i via a stream r (e.g., **yield** array[t_i]). Fig. 18 illustrates abstract ABS code for *requests* which streams the values, and for *delegates* which receives them.

7.3. Experimental results

The case study on massive social network simulation has been implemented in Cloud ABS [7], which is a source-to-source trans-compiler from ABS code down to Cloud Haskell [18] runnable on distributed machines. Beside a higher level of abstraction at the programming level thanks to our proposed feature, the distributed run-time system provides more than $6\times$ speed-up in performance compared to the same implementation without the use of the feature, as presented in [12]. The results are illustrated in Fig. 19.

The distribution overhead increases the execution time for two machines, which is compensated by the parallelism achieved through adding more VMs. As shown in Table 1, the memory consumption decreases when adding more VMs, which enables the generation of extra-large graphs which don't fit in centralized-memory architectures. We ran the implementation on a distributed cloud environment kindly provided by the Dutch SURF foundation. The hardware consisted of




Fig. 19. Performance results of the distributed PA in ABS-Haskell for graphs of $n = 10^7$ nodes with $d =$ (a) 3, (b) 10.

Table 1

Maximum memory residency (in MB) per Virtual Machine.

Graph size	Total number of VMs				
	1	2	4	8	16
$n = 10^6, d = 3$	306	266	212	155	114
$n = 10^6, d = 10$	899	1028	547	354	221
$n = 10^7, d = 3$	2123	3242	1603	967	621
$n = 10^7, d = 10$	6260	9668	6702	3611	1905

identical VMs interconnected over a 10 Gbps Ethernet network. Each VM was a single-core Intel Xeon E5-2698, 16 GB RAM running Ubuntu 14.04 Server edition. We provided an online repository³ containing the ABS code for the case study. Also the instructions for installing the ABS Haskell backend are provided online.⁴

8. Related work

There already exists a variety of different programming constructs for streaming data in different programming languages like Java, and software frameworks for processing big data like Apache Hadoop and Spark.

Asynchronous generators specified in [19] enable the streaming of data in an asynchronous method invocation. This includes, on the callee side, yielding the data, and on the caller side receiving them as an asynchronous iterator or raising an exception if there is no further yielded data. These generators are defined in the context of the multi-threaded model of concurrency, where asynchrony is provided by spawning a thread for a method call.

Akka Streams [20] provides an API to specify a streaming setup between actors which allows to adapt behavior to the underlying resources in terms of both memory and speed.

There are also languages which utilize the notion of *channel* as a means of communication, inspired by the model of *Communicating Sequential Processes* (CSP). For instance, Go language and *JCSP* [21], which is a library in Java, provide CSP-like elements, e.g., processes (referred to as Goroutines in Go) that communicate via channels by means of read and write primitives. Buffered channels in Go provide asynchronous read (cf. write) when the buffer is not empty (cf. not full). Otherwise the primitives are blocking.

Similarly to asynchronous generators, streaming data as proposed in this paper is fully integrated with asynchronous method invocation, i.e., it is not a separate orthogonal concept like channels are. But its integration with the ABS language allows for an additional loose coupling between the producer and consumer of data streams: by means of cooperative scheduling of tasks the consumption of data can be interleaved with other tasks on demand.

The distributed shared memory (DSM) paradigm [22,23,10] enables access to a common shared space across disjoint address spaces, where communication and synchronization between processes are enforced through operations on shared data. The notion of *tuple space* was originally integrated at the language level in Linda [24]. The processes communicate via insertion/read/removal of tuples into/from the tuple: `out()` to write a tuple into a tuple space, `in()` to retrieve (and remove), and `read()` to read (without removing) a tuple from it.

Similarly to tuple spaces, the interaction model of streams proposed in this paper provides time and space decoupling, namely, data producers and consumers can remain anonymous with respect to each other, and the sender of a data needs no knowledge about the future use of that data or its destination (the reference to a stream can be passed around). Also producer-side synchronization decoupling is guaranteed, whereas, the consumer-side decoupling is not provided in tuple-spaces, as the consumers synchronously pull the data. In ABS streams, however, the decoupling is provided at the consumer

³ <https://github.com/kazadbakht/PA/blob/master/src/packaging/DisPAck.abs>.

⁴ <https://github.com/abstools/habs>.

object level, thanks to the notion of cooperative scheduling. Similarly to tuple space in-based and read-based communication, the destructive and non-destructive data streams, respectively, can be naturally used to implement *one-of-n* semantics (only one consumer reads a given data), and *one-to-n* message delivery (a given data can be read by all such consumers).

9. Future work

We focused on extending the main asynchronous core of ABS with data streams. Other main features of the ABS like *concurrent object groups* (cogs) and deployment components are orthogonal and compatible with this extension. As an example, ABS features *cog* that, in principle, shares a thread of control among its constituent objects, which enables internal synchronous calls. By the nature of data streaming, it is natural to restrict the streaming to the asynchronous method calls.

In this paper we discussed and used an implementation of the ABS model of data streams in the Haskell backend. The ABS with Haskell backend supports real-time programming techniques which allows for specifying deadlines with method invocations. This provides an interesting basis to extend ABS with real-time data streaming which may, as an example, involve timeout on read operations. We also need to extend the various formal analysis techniques (e.g., deadlock detection, general functional analysis based on method contracts) currently supported by the ABS to the ABS model of streaming data discussed in this paper.

Acknowledgements

The experiments described in this paper were carried out on the Dutch national HPC cloud infrastructure, a service provided by the SURF Foundation (<http://www.surf.nl>). We thank anonymous reviewers for their constructive comments that helped improve and clarify the manuscript.

References

- [1] G.A. Agha, *Actors: a Model of Concurrent Computation in Distributed Systems*, Technical Report, Massachusetts Inst of Tech Cambridge Artificial Intelligence Lab, 1985.
- [2] J. Armstrong, R. Virding, C. Wikström, M. Williams, *Concurrent Programming in Erlang*, 1993.
- [3] E.B. Johnsen, R. Hähnle, J. Schäfer, R. Schlatte, M. Steffen, ABS: a core language for abstract behavioral specification, in: *International Symposium on Formal Methods for Components and Objects*, Springer, 2010, pp. 142–164.
- [4] F.S. De Boer, D. Clarke, E.B. Johnsen, A complete guide to the future, in: *European Symposium on Programming*, Springer, 2007, pp. 316–330.
- [5] M. Sirjani, Rebeca: theory, applications, and tools, in: *Formal Methods for Components and Objects*, 2006, pp. 102–126.
- [6] K. Azadbakht, N. Bezirgiannis, F.S. de Boer, On futures for streaming data in ABS, in: *International Conference on Formal Techniques for Distributed Objects, Components, and Systems*, Springer, 2017, pp. 67–73.
- [7] N. Bezirgiannis, F. de Boer, ABS: a high-level modeling language for cloud-aware programming, in: *International Conference on Current Trends in Theory and Practice of Computer Science (SOFSEM)*, Springer, 2016, pp. 433–444.
- [8] E. Giachino, C. Laneve, M. Lienhardt, A framework for deadlock detection in core ABS, *Softw. Syst. Model.* 15 (2016) 1013–1048.
- [9] K. Azadbakht, F.S. de Boer, E. de Vink, Deadlock detection for actor-based coroutines, in: *International Symposium on Formal Methods*, Springer, 2018, pp. 39–54.
- [10] P.T. Eugster, P.A. Felber, R. Guerraoui, A.-M. Kermarrec, The many faces of publish/subscribe, *ACM Comput. Surv.* 35 (2003) 114–131.
- [11] G.D. Plotkin, *A Structural Approach to Operational Semantics*, 1981.
- [12] K. Azadbakht, N. Bezirgiannis, F.S. de Boer, Distributed network generation based on preferential attachment in ABS, in: *International Conference on Current Trends in Theory and Practice of Informatics*, Springer, 2017, pp. 103–115.
- [13] J. Leskovec, *Dynamics of Large Networks*, ProQuest, 2008.
- [14] D. Bader, K. Madduri, et al., Parallel algorithms for evaluating centrality indices in real-world networks, in: *Parallel Processing, ICPP 2006. International Conference on*, IEEE, 2006, pp. 539–550.
- [15] A.-L. Barabási, R. Albert, Emergence of scaling in random networks, *Science* 286 (1999) 509–512.
- [16] K. Azadbakht, N. Bezirgiannis, F.S. de Boer, S. Aliakbary, A high-level and scalable approach for generating scale-free graphs using active objects, in: *Proceedings of the 31st Annual ACM Symposium on Applied Computing*, ACM, 2016, pp. 1244–1250.
- [17] R. Kumar, P. Raghavan, S. Rajagopalan, D. Sivakumar, A. Tomkins, E. Upfal, Stochastic models for the web graph, in: *Foundations of Computer Science, 2000. Proceedings. 41st Annual Symposium on*, IEEE, 2000, pp. 57–65.
- [18] J. Epstein, A.P. Black, S. Peyton-Jones, Towards Haskell in the Cloud, *ACM SIGPLAN Notices*, vol. 46, ACM, 2011, pp. 118–129.
- [19] Y. Selivanov, Asynchronous generators, <https://www.python.org/dev/peps/pep-0525/>, 2016.
- [20] Streams, Version 2.4.17, <http://doc.akka.io/docs/akka/2.4/scala/stream/index.html>, 2017.
- [21] P. Welch, N. Brown, Communicating sequential processes for java[™] (jcs), <https://www.cs.kent.ac.uk/projects/ofa/jcsp/>, 2014.
- [22] K. Li, P. Hudak, Memory coherence in shared virtual memory systems, *ACM Trans. Comput. Syst.* 7 (1989) 321–359.
- [23] M.-C. Tam, J.M. Smith, D.J. Farber, A taxonomy-based comparison of several distributed shared memory systems, *ACM SIGOPS Oper. Syst. Rev.* 24 (1990) 40–67.
- [24] D. Gelernter, Generative communication in linda, *ACM Trans. Program. Lang. Syst.* 7 (1985) 80–112.