

# An Executable Sequential Specification for Spark Aggregation

Yu-Fang Chen<sup>1</sup>, Chih-Duo Hong<sup>1</sup>, Ondřej Lengál<sup>1,2</sup>,  
Shin-Cheng Mu<sup>1</sup>, Nishant Sinha<sup>3</sup>, Bow-Yaw Wang<sup>1</sup>

<sup>1</sup> Academia Sinica, Taiwan

<sup>2</sup> Brno University of Technology, Czech Republic

<sup>3</sup> IBM Research, India

**Abstract.** Spark is a new promising platform for scalable data-parallel computation. It provides several high-level application programming interfaces (APIs) to perform parallel data aggregation. Since execution of parallel aggregation in Spark is inherently non-deterministic, a natural requirement for Spark programs is to give the same result for any execution on the same data set. We present PURESPARK, an executable formal Haskell specification for Spark aggregate combinators. Our specification allows us to deduce the precise condition for deterministic outcomes from Spark aggregation. We report case studies analyzing deterministic outcomes and correctness of Spark programs.

## 1 Introduction

Spark [30,1,31] is a popular platform for scalable distributed data-parallel computation based on a flexible programming environment with concise and high-level APIs. Spark is by many considered as the successor of MapReduce [14,26]. Despite its fame, the precursory computational model of MapReduce suffers from I/O congestion and limited programming support for distributed problem solving. Notably, Spark has the following advantages over MapReduce. First, it has high performance due to distributed, cached, and in-memory computation. Second, the platform adopts a relaxed fault tolerant model where sub-results are recomputed upon faults rather than aggressively stored. Third, lazy evaluation semantics is used to avoid unnecessary computation. Finally, Spark offers greater programming flexibility through its powerful APIs founded in functional programming. Spark also owes its popularity to a unified framework for efficient graph, streaming, and SQL-based relational database computation, a machine learning library, and the support of multiple distributed data storage formats. Spark is one of the most active open-source projects with over 1000 contributors [1].

In a typical Spark program, a sequence of transformations followed by an action are performed on Resilient Distributed Datasets (RDDs). An RDD is the principal abstraction for data-parallel computation in Spark. It represents a read-only collection of data items partitioned and stored distributively. RDD operations such as `map`, `reduce`, and `aggregate` are called *combinators*. They generate and aggregate data in RDDs to carry out Spark computation. For instance, the `aggregate` combinator takes user-defined functions `seq` and `comb`: `seq` accumulates a sub-result for each partition while `comb` merges sub-results across different partitions. Spark also provides a family of aggregate combinators for common data structures such as pairs and graphs. In Spark computation, data aggregation is ubiquitous.

Programming in Spark, however, can be tricky. Since sub-results are computed using multiple applications of *seq* and *comb* across partitions concurrently, the order of their applications varies on different executions. Because of indefinite orders of computation, aggregation in Spark is inherently *non-deterministic*. A Spark program may produce different outcomes for the same input on different runs. This form of non-deterministic computation has other side effects. For instance, the private function `AreaUnderCurve.of` in the Spark machine learning library computes numerical integration distributively; it exhibits numerical instability due to non-deterministic computation. Consider the integral of  $x^{73}$  on the interval  $[-2, 2]$ . Since  $x^{73}$  is an odd function, the integral is 0. In our experiments, `AreaUnderCurve.of` returns different results ranging from  $-8192.0$  to  $12288.0$  on the same input because of different orders of floating-point computation. To ensure deterministic outcomes, programmers must carefully develop their programs to adhere to Spark requirements.

Unfortunately, Spark’s documentation does not specify the requirements formally. It only describes informal algebraic properties about combinators to ensure correctness. The documentation provides little help to a programmer in understanding the complex, and sometimes unexpected, interaction between *seq* and *comb*, especially when these two are functions over more complex domains, e.g. lists or trees. Inspecting the Spark implementation is a laborious job since public combinators are built by composing a long chain of generic private combinators—determining the execution semantics from the complex implementation is hard. Moreover, Spark is continuously evolving and the implementation semantics may change significantly across releases. We therefore believe that a formal specification of Spark combinators is necessary to help developers understand the program semantics better, clarify hidden assumptions about RDDs, and help to reason about correctness and sources of non-determinism in Spark programs.

Building a formal specification for Spark is far from straightforward. Spark is implemented in Scala and provides high-level APIs also in Python and Java. Because Spark heavily exploits various language features of Scala, it is hard to derive specifications without formalizing the operational semantics of the Scala language, which is not an easy task by itself. Instead of that, we have developed a Haskell library `PURESPARK` [4], which for each key Spark combinator provides an abstract sequential functional specification in Haskell. We use Haskell as a specification language for two reasons. First, the core of Haskell has strong formal foundations in  $\lambda$ -calculus. Second, program evaluation in Haskell, like Spark, is lazy, which admits faithful modeling of Spark aggregation. Through the use of Haskell we obtain a concise formal functional model for Spark combinators without formalizing Scala.

An important goal of our specification is to make non-determinism in various combinators explicit. Spark developers can inspect it to identify sources of non-determinism when program executions yield unexpected outputs. Researchers can also use it to understand distributed Spark aggregation and investigate its computational pattern. Our specification is also *executable*. A programmer can use the Haskell APIs to implement data-parallel programs, test them on different input RDDs, and verify correctness of outputs independent of the Spark programming environment. In our case studies, we capture non-deterministic behaviors of real Spark programs by executing the corresponding `PURESPARK` specifications with crafted input data sets. We also show that the sequential specification is useful in developing distributed Spark programs.

Our main contributions are summarized below:

- We present formal, functional, sequential specifications for key Spark aggregate combinators. The PURESPARK specification consists of executable library APIs. It can assist Spark program development by mimicking data-parallel programming in conventional environments.
- Based on the specification, we investigate and identify necessary and sufficient conditions for Spark aggregate combinators to produce deterministic outcomes for general and pair RDDs.
- Our specification allows to deduce the precise condition for deterministic outcomes from Spark aggregation.
- We perform a series of case studies on practical Spark programs to validate our formalization. With PURESPARK, we find instances of numerical instability in the Spark machine learning library.
- Up to our knowledge, this is the first work to provide a formal, functional specification of key Spark aggregate combinators for data-parallel computation.

## 2 Preliminaries

Let  $A$  be a non-empty set and  $\odot : A \times A \rightarrow A$  be a function. An element  $i \in A$  is the *identity* of  $\odot$  if for every  $a \in A$ , it holds that  $a = i \odot a = a \odot i$ . The function  $\odot$  is *associative* if for every  $a, a', a'' \in A$ ,  $a \odot (a' \odot a'') = (a \odot a') \odot a''$ ;  $\odot$  is *commutative* if for every  $a, a' \in A$ ,  $a \odot a' = a' \odot a$ . The algebraic structure  $(A, \odot)$  is a *semigroup* if  $\odot$  is associative. A *monoid* is a structure  $(A, \odot, \perp)$  such that  $(A, \odot)$  is a semigroup and  $\perp \in A$  is the identity of  $\odot$ . The semigroup  $(A, \odot)$  and monoid  $(A, \odot, \perp)$  are commutative if  $\odot$  is commutative.

Haskell is a strongly typed purely functional programming language. Similar to Spark, Haskell programs are lazily evaluated. We use several widely used Haskell functions (Figure 1). `fst` and `snd` are projections on pairs. `null` tests whether a list is empty. `elem` is the membership function for lists; its infix notation is often used, as in `0 `elem` []`. `(++)` concatenates two lists; it is used as an infix operator, as in `[False] ++ [True]`. `map` applies a function to elements of a list. `reduce1` merges elements of a list by a given binary function from left to right. `fold1` accumulates by applying a function to elements of a list iteratively, also from left to right. `concat` concatenates elements in a list. `concatMap` applies a function to elements of a list and concatenates the results. `lookup` finds the value of a key in a list of pairs. `filter` selects elements from a list by a predicate.

In order to formalize non-determinism in distributed aggregation, we define the following non-deterministic shuffle function for lists:

```
shuffle! :: [ $\alpha$ ]  $\rightarrow$  [ $\alpha$ ]
shuffle! xs = ...      -- shuffle xs randomly
```

A random monad can be used to define random shuffling. Instead of explicit monadic notation, we introduce the *chaotic* `shuffle!` function in our presentation for the sake of brevity. Thus, `shuffle! [0, 1, 2]` evaluates to one of the six possible lists `[0, 1, 2]`, `[0, 2, 1]`, `[1, 0, 2]`, `[1, 2, 0]`, `[2, 0, 1]`, or `[2, 1, 0]` randomly. Using `shuffle!`, more chaotic functions are defined.

```
map! :: ( $\alpha \rightarrow \beta$ )  $\rightarrow$  [ $\alpha$ ]  $\rightarrow$  [ $\beta$ ]
map! f xs = shuffle! (map f xs)

concatMap! :: ( $\alpha \rightarrow [\beta]$ )  $\rightarrow$  [ $\alpha$ ]  $\rightarrow$  [ $\beta$ ]
concatMap! f xs = concat (map! f xs)
```



More concretely, let  $z$  be a default aggregated value. `aggregate` applies `foldl seq z` to every partition of `rdd`. Hence the sub-result of each partition is accumulated by folding elements in the partition with `seq`. The combinator then combines sub-results by another folding using `comb`.

Note that the chaotic `map!` function is used to model non-deterministic interleavings of sub-results. To exploit concurrency, Spark creates a task to compute the sub-result for each partition. These tasks are executed concurrently and hence induce non-deterministic computation. We use the chaotic `map!` function to designate non-determinism explicitly.

A related combinator is `reduce`. Instead of `foldl`, the combinator uses `reducel` to aggregate data in an RDD.

```
reduce :: ( $\alpha \rightarrow \alpha \rightarrow \alpha$ )  $\rightarrow$  RDD  $\alpha \rightarrow \alpha$ 
reduce comb rdd = let presults = map! (reducel comb) rdd
                  in reducel comb presults
```

Similar to the `aggregate` combinator, `reduce` computes sub-results concurrently. The chaotic `map!` function is again used to model non-deterministic computation.

Sub-results of different partitions are computed in parallel, but the `aggregate` combinator still combines sub-results sequentially. This can be further parallelized. Observe that several sub-results may be available simultaneously from distributed computation. The Spark `treeAggregate` combinator applies `comb` to pairs of sub-results concurrently until the final result is obtained. In addition to concurrent computation of sub-results, `treeAggregate` also combines sub-results from different partitions in parallel.

In our specification, two chaotic functions are used to model non-deterministic computation on two different levels. The `map!` function models non-determinism in computing sub-results of partitions. The `apply!` function (introduced below) models concurrent combination of sub-results from different partitions. It combines two consecutive sub-results picked chaotically, and repeats such chaotic combinations until the final result is obtained. Observe that the computation has a binary-tree structure with `comb` as internal nodes and sub-results from different partitions as leaves.

```
apply! :: ( $\beta \rightarrow \beta \rightarrow \beta$ )  $\rightarrow$  [ $\beta$ ]  $\rightarrow \beta$ 
apply! comb [r] = r
apply! comb [r, r'] = comb r r'
apply! comb rs = let (ls', l', r', rs') = ... -- rs == ls' ++ [l', r'] ++ rs'
                  in apply! comb (ls' ++ [comb l' r'] ++ rs')
```

```
treeAggregate ::  $\beta \rightarrow (\beta \rightarrow \alpha \rightarrow \beta) \rightarrow (\beta \rightarrow \beta \rightarrow \beta) \rightarrow$  RDD  $\alpha \rightarrow \beta$ 
treeAggregate z seq comb rdd = let presults = map! (foldl seq z) rdd
                               in apply! comb presults
```

The `treeReduce` combinator optimizes `reduce` by combining sub-results in parallel. Similar to `treeAggregate`, two levels of non-deterministic computation can occur.

```
treeReduce :: ( $\alpha \rightarrow \alpha \rightarrow \alpha$ )  $\rightarrow$  RDD  $\alpha \rightarrow \alpha$ 
treeReduce comb rdd = let presults = map! (reducel comb) rdd
                      in apply! comb presults
```

*Pair RDDs.* Key-value pairs are widely used in data parallel computation. If the data type of an RDD is a pair, we say that the RDD is a *pair* RDD. The first and second elements in a pair are called the *key* and the *value* of the pair respectively.

**type** PairRDD  $\alpha \beta = \text{RDD } (\alpha, \beta)$

In a pair RDD, different pairs can have the same key. Spark provides combinators to aggregate values associated with the same key. The `aggregateByKey` combinator returns an RDD by aggregating values associated with the same key. We use the following functions to formalize `aggregateByKey`:

```

hasKey ::  $\alpha \rightarrow \text{Partition } (\alpha, \beta) \rightarrow \text{Bool}$ 
hasKey k ps = case (lookup k ps) of
  Just _  $\rightarrow \text{True}$ 
  Nothing  $\rightarrow \text{False}$ 

hasValue ::  $\alpha \rightarrow \beta \rightarrow \text{Partition } (\alpha, \beta) \rightarrow \beta$ 
hasValue k val ps = case (lookup k ps) of
  Just v  $\rightarrow v$ 
  Nothing  $\rightarrow \text{val}$ 

```

```

addTo ::  $\alpha \rightarrow \beta \rightarrow \text{Partition } (\alpha, \beta) \rightarrow \text{Partition } (\alpha, \beta)$ 
addTo key val ps = foldl ( $\lambda r (k, v) \rightarrow \text{if } \text{key} == k \text{ then } r \text{ else } (k, v):r$ ) [(key, val)] ps

```

The expression `hasKey key ps` checks if `key` appears in a partition of pairs. `hasValue key val ps` finds a value associated with `key` in a partition of pairs. It evaluates to the default value `val` if `key` does not appear in the partition. The expression `addTo key val ps` adds the pair `(key, val)` to the partition `ps`, and removes other pairs with the same key.

The `aggregateByKey` combinator first aggregates all pairs with the value `z` and the function `mergeComb` in each partition. If values `vs` are associated with the same key in a partition, the value `foldl mergeComb z vs` for the key is pre-aggregated. Since a key may appear in several partitions, all pre-aggregated values associated with the key across different partitions are merged using `mergeValue`.

```

aggregateByKey ::  $\gamma \rightarrow (\gamma \rightarrow \beta \rightarrow \gamma) \rightarrow (\gamma \rightarrow \gamma \rightarrow \gamma) \rightarrow \text{PairRDD } \alpha \beta \rightarrow \text{PairRDD } \alpha \gamma$ 
aggregateByKey z mergeComb mergeValue pairRdd =
  let mergeBy fun left (k, v) = addTo k (fun (hasValue k z left) v) left
    preAgg = concatMap! (foldl (mergeBy mergeComb) []) pairRdd
  in repartition! (foldl (mergeBy mergeValue) [] preAgg)

```

In the specification, we accumulate values associated with the same key by `mergeComb` in each partition, keeping a list of pairs of a key and the partially aggregated value for the key. Since accumulation in different partitions runs in parallel, the chaotic `concatMap!` function is used to model such non-deterministic computation. After all partitions finish their accumulation, `mergeValue` merges values associated with the same key across different partitions. The final pair RDD can have a default or user-defined partitioning. Since a user-defined partitioning may shuffle a pair RDD arbitrarily, it is in our specification modeled by the chaotic `repartition!` function.

Pair RDDs have a combinator corresponding to reduce called `reduceByKey`. The `reduceByKey` function merges all values associated with a key by `mergeValue`. It follows a similar computational pattern as `aggregateByKey`. Note that every key is associated with at most one value in resultant pair RDDs of `aggregateByKey` or `reduceByKey`.

```

reduceByKey ::  $(\beta \rightarrow \beta \rightarrow \beta) \rightarrow \text{PairRDD } \alpha \beta \rightarrow \text{PairRDD } \alpha \beta$ 
reduceByKey mergeValue pairRdd =
  let merge left (k, v) = case lookup k left of Just v'  $\rightarrow \text{addTo } k \text{ (mergeValue } v' \text{ v) left}$ 
    Nothing  $\rightarrow \text{addTo } k \text{ v left}$ 
    preAgg = concatMap! (foldl merge []) pairRdd
  in repartition! (foldl merge [] preAgg)

```

Spark also provides a library, called GraphX, for a distributed analysis of graphs. See App A for a formalization of some of its key functions.

## 4 Deterministic Aggregation

Having deterministic outcomes is desired from all aggregation functions. If a function may return different values on different executions, the function is often not implemented correctly. A program with explicit assumptions on the input data is also desirable. Otherwise, the program may work correctly on certain data sets but produce unexpected outcomes on others where implicit assumptions do not hold [28]. We now investigate conditions under which Spark aggregation combinators always produce deterministic outcomes. Proofs of the given lemmas can be found in App. C. Proofs of some crucial lemmas have also been formalized using Agda [4].

We first show how to deal with non-deterministic behaviors in the aggregate combinator. Consider a variant of the formalization of aggregate from Section 3:

```
aggregate' :: β → (β → α → β) → (β → β → β) → RDD α → β
aggregate' z seq comb rdd = let results = perm (map (foldl seq z) rdd)
  in foldl comb z results
```

Observe that we changed the application of the chaotic map! function with an application of the permutation perm after the regular map function. The function composition perm( map ...) is a concrete instantiation of map!, that is, a function that permutes its list argument. Notice that perm can be pushed inside map:

```
perm (map f xs) == map f (perm xs).
```

Assume that rdd was obtained from a list xs by splitting and permuting, that is, rdd == perm' (split xs) where split :: [α] → [[α]] satisfies xs == (concat . split) xs. We can therefore rewrite the computation of results in aggregate' to

```
let pres = perm (map (foldl seq z) (perm' (split xs))),
```

After pushing perm inside map, we obtain

```
let pres = map (foldl seq z) ((perm . perm') (split xs)).
```

Since perm . perm' is also a permutation perm'', we have

```
let pres = map (foldl seq z) rdd'
```

where rdd' is another RDD obtained from xs by splitting and shuffling. Let us call (deterministic) instances of repartition! as *partitionings*. As a consequence, we focus only on proving if calls to aggregate<sup>D</sup> defined below have deterministic outcomes for different partitionings of a list into RDDs:

```
aggregateD :: β → (β → α → β) → (β → β → β) → RDD α → β
aggregateD z seq comb rdd = let pres = map (foldl seq z) rdd
  in foldl comb z pres
```

-3oreover, we define a deterministic versions of reduce:

```
reduceD :: (α → α → α) → RDD α → α
reduceD comb rdd = let results = perm (map (reducel comb) rdd)
  in reducel comb results
```

and also treeAggregate<sup>D</sup> and treeReduce<sup>D</sup> in a similar way.

In the following, given a function f that takes an RDD as one of its parameters and contains a single occurrence of the chaotic map! (respectively concatMap!) function, we use f<sup>D</sup> to denote the function obtained from f by replacing the chaotic map! (respectively concatMap!) with a regular map (respectively concatMap). A similar reasoning

can show that it suffices to check whether calls to  $f^D$  have deterministic outcomes for different partitionings on a list into RDDs.

For better readability, standard mathematical notation of functions is used in the rest of this section. We represent a Haskell function application  $f\ x_1 \dots x_n$  as  $f(x_1, \dots, x_n)$ .

#### 4.1 aggregate

In this section, we give conditions for deterministic outcomes of calls to the aggregate combinator  $\text{aggregate}(z, \text{seq}, \oplus, \text{rdd})$  for  $z :: \beta$ ,  $\text{seq} :: \beta \times \alpha \rightarrow \beta$ ,  $\oplus :: \beta \times \beta \rightarrow \beta$ , and  $\text{rdd} :: \text{RDD } \alpha$ . We first define what it means for calls to the aggregate combinator to have deterministic outcomes.

**Definition 1.** *Calls to  $\text{aggregate}(z, \text{seq}, \oplus, \text{rdd})$  have deterministic outcomes if*

$$\text{aggregate}^D(z, \text{seq}, \oplus, \text{part}(L)) = \text{foldl}(\text{seq}, z, L) \quad (1)$$

for all lists  $L$  and partitionings  $\text{part}$ .

Conventionally,  $\text{aggregate}$  is regarded as a parallelized counterpart of  $\text{foldl}$ . For example, the sequential  $\text{aggregate}$  function in the standard Scala library ignores the  $\oplus$  operator and is implemented by  $\text{foldl}$ . This is why we characterize deterministic  $\text{aggregate}$  as  $\text{foldl}$  in Definition 1. Our characterization, however, does not cover all  $\text{aggregate}$  calls that always give the same outputs. In particular, it does not cover an  $\text{aggregate}$  call where  $\oplus$  is a constant function, which is, however, quite suspicious in a distributed data-parallel computation and should be reported.

We give necessary and sufficient conditions for  $\text{aggregate}$  calls to have deterministic outcomes in several lemmas, culminating in Corollary 1. The first lemma allows us to check only conditions on  $\text{seq}$  and  $\oplus$  over all possible pairs of lists instead of enumerating all possible partitionings on lists. For brevity, we use  $\langle p_1 \rangle$  for  $\text{foldl}(\text{seq}, z, p_1)$ , and  $\text{img}(\text{foldl}(\text{seq}, z))$  for the image of  $\text{foldl}(\text{seq}, z, L)$  for any list  $L$ . That is,  $\text{img}(\text{foldl}(\text{seq}, z)) = \{y \mid \text{there is a list } L \text{ such that } \text{foldl}(\text{seq}, z, L) = y\}$ .

**Lemma 1.** *Calls to  $\text{aggregate}(z, \text{seq}, \oplus, \text{rdd})$  have deterministic outcomes iff:*

1. *( $\text{img}(\text{foldl}(\text{seq}, z)), \oplus, z$ ) is a commutative monoid, and*
2. *for all lists  $p_1, p_2 :: [\alpha]$ ,  $\langle p_1 \uplus p_2 \rangle = \langle p_1 \rangle \oplus \langle p_2 \rangle$ .*

Note that condition 2 in Lemma 1 is equivalent to saying that  $\langle \cdot \rangle$  is a list homomorphism to the monoid  $(\text{img}(\text{foldl}(\text{seq}, z)), \oplus, z)$  [6].

The lemma below further helps us reduce the need of testing conditions over all possible pairs of lists to conditions over elements of  $\alpha \times \text{img}(\text{foldl}(\text{seq}, z))$ .

**Lemma 2.** *Let  $\oplus$  be associative on  $\gamma = \text{img}(\text{foldl}(\text{seq}, z))$  and  $z$  be the identity of  $\oplus$  on  $\gamma$ . The following are equivalent:*

1. *for all lists  $p_1, p_2 :: [\alpha]$ ,*

$$\langle p_1 \uplus p_2 \rangle = \langle p_1 \rangle \oplus \langle p_2 \rangle, \quad (2)$$

2. *for all elements  $d :: \alpha$  and  $e :: \gamma$ ,*

$$\text{seq}(e, d) = e \oplus \text{seq}(z, d). \quad (3)$$

Summarizing the lemmas, we get the following corollary:

**Corollary 1.** *Calls to  $\text{aggregate}(z, \text{seq}, \oplus, \text{rdd})$  have deterministic outcomes iff*

1. *( $\text{img}(\text{foldl}(\text{seq}, z)), \oplus, z$ ) is a commutative monoid and*
2. *for all  $d :: \alpha$  and  $e :: \text{img}(\text{foldl}(\text{seq}, z))$ , it holds that  $\text{seq}(e, d) = e \oplus \text{seq}(z, d)$ .*

## 4.2 reduce

This section explores conditions for deterministic outcomes of calls to  $\text{reduce}(\oplus, rdd)$  for  $\oplus :: \alpha \times \alpha \rightarrow \alpha$  and  $rdd :: \text{RDD } \alpha$ . We use the function  $\text{reduce}^D$  defined in the introduction of Section 4. For  $\text{reduce}$ , we assume that for any non-empty list, all partitions of its partitioning are non-empty (otherwise the result of  $\text{reduce}$  is undefined).

We define deterministic outcomes for  $\text{reduce}$  as follows.

**Definition 2.** *Calls to  $\text{reduce}(\oplus, rdd)$  have deterministic outcomes if*

$$\text{reduce}^D(\oplus, \text{part}(L)) = \text{reduce}(\oplus, L) \quad (4)$$

for all lists  $L$  and partitionings  $\text{part}$ .

We reduce the problem of checking if  $\text{reduce}$  has deterministic outcomes to the problem of checking if  $\text{aggregate}$  has deterministic outcomes by the following lemma.

**Lemma 3.** *Calls to  $\text{reduce}(\oplus, rdd)$  have deterministic outcomes iff calls to  $\text{aggregate}(\text{Nothing}, \text{seq}', \oplus', rdd)$  have deterministic outcomes, where  $\text{seq}'$  and  $\oplus'$  are as follows:*

$$\begin{array}{ll} \text{seq}' \times y = \text{case } x \text{ of} & (\oplus') \times y = \text{case } (x, y) \text{ of } (\text{Nothing}, y') \rightarrow y' \\ \text{Nothing} \rightarrow \text{Just } y & (x', \text{Nothing}) \rightarrow x' \\ \text{Just } x' \rightarrow \text{Just } (x' \oplus y) & (\text{Just } x', \text{Just } y') \rightarrow \text{Just } (x' \oplus y'). \end{array}$$

Combining Corollary 1 and Lemma 3, we get the condition for deterministic outcomes of  $\text{reduce}(\oplus, rdd)$  calls.

**Corollary 2.** *Calls to  $\text{reduce}(\oplus, rdd)$  have deterministic outcomes iff  $(\alpha, \oplus)$  is a commutative semigroup.*

## 4.3 treeAggregate and treeReduce

This section gives conditions for deterministic outcomes of calls to the following two aggregate combinators:

1.  $\text{treeAggregate}(z, \text{seq}, \oplus, rdd)$  for  $z :: \beta$ ,  $\text{seq} :: \beta \times \alpha \rightarrow \beta$ ,  $\oplus :: \beta \times \beta \rightarrow \beta$ , and  $rdd :: \text{RDD } \alpha$ ; and
2.  $\text{treeReduce}(\oplus, rdd)$  for  $\oplus :: \alpha \times \alpha \rightarrow \alpha$ ,  $rdd :: \text{RDD } \alpha$ .

Different from  $\text{aggregate}$  and  $\text{reduce}$ , the tree variants have another level of non-determinism modeled by  $\text{apply!}$ . The chaotic function effectively simulates non-deterministic computation with a binary-tree structure (Section 3).

To define calls to  $\text{treeAggregate}$  and  $\text{treeReduce}$  to have deterministic outcomes, we use the functions  $\text{treeAggregate}^T$  and  $\text{treeReduce}^T$  obtained by adding an explicit deterministic instantiation of  $\text{apply!}$  to  $\text{treeAggregate}^D$  and  $\text{treeReduce}^D$ .

**Definition 3.** *Calls to  $\text{treeAggregate}(z, \text{seq}, \oplus, rdd)$  and  $\text{treeReduce}(\oplus, rdd)$  have deterministic outcomes if*

$$\text{treeAggregate}^T(\text{apply}, z, \text{seq}, \oplus, \text{part}(L)) = \text{foldl}(\text{seq}, z, L) \quad (5)$$

and

$$\text{treeReduce}^T(\text{apply}, \oplus, \text{part}(L)) = \text{reduce}(\oplus, L) \quad (6)$$

respectively for all lists  $L$ , partitionings  $\text{part}$ , and instantiations  $\text{apply}$  of  $\text{apply!}$ .

The following two propositions state necessary and sufficient conditions for the  $\text{treeAggregate}$  and  $\text{treeReduce}$  combinators to have deterministic outcomes.

**Proposition 1.** *Calls to  $\text{treeAggregate}(z, \text{seq}, \oplus, rdd)$  have deterministic outcomes iff calls to  $\text{aggregate}(z, \text{seq}, \oplus, rdd)$  have deterministic outcomes.*

**Proposition 2.** *Calls to  $\text{treeReduce}(\oplus, rdd)$  have deterministic outcomes iff calls to  $\text{reduce}(\oplus, rdd)$  have deterministic outcomes.*

#### 4.4 aggregateByKey and reduceByKey

We proceed by investigating conditions for the following combinators on pair RDDs:

1. `aggregateByKey(z, seq, ⊕, prdd)` for  $z :: \gamma$ ,  $seq :: \gamma \times \beta \rightarrow \gamma$ ,  $\oplus :: \gamma \times \gamma \rightarrow \gamma$ , and  $prdd :: \text{PairRDD } \alpha \beta$ ; and
2. `reduceByKey(⊕, prdd)` for  $\oplus :: \beta \times \beta \rightarrow \beta$  and  $prdd :: \text{PairRDD } \alpha \beta$ .

We define an auxiliary function `filterkey` that obtains a list of all values associated with the given key from a list of pairs.

**filterkey**  $:: \alpha \rightarrow [(\alpha, \beta)] \rightarrow [\beta]$

**filterkey**  $- [] = []$

**filterkey**  $k (k, v):xs = v:(\text{filterkey } k \text{ xs})$

**filterkey**  $k (_, _):xs = \text{filterkey } k \text{ xs}$

Deterministic outcomes of calls to `aggregateByKey` are now defined using the function `aggregateByKeyD` as follows.

**Definition 4.** *Calls to `aggregateByKey(z, seq, ⊕, prdd)` have deterministic outcomes if*

$$\text{lookup}(k, \text{aggregateByKey}^D(z, seq, \oplus, \text{part}(L))) = \text{foldl}(z, seq, \text{filterkey}(k, L))$$

for all lists  $L$  of pairs, partitionings  $\text{part}$ , and keys  $k$ .

Finally, the following proposition states the conditions that need to hold for calls to `aggregateByKey` to have deterministic outcomes.

**Proposition 3.** *Calls to `aggregateByKey(z, seq, ⊕, prdd)` have deterministic outcomes iff calls to `aggregate(z, seq, ⊕, rdd)` have deterministic outcomes.*

We define when calls to `reduceByKey` have deterministic outcomes via `reduceByKeyD`.

**Definition 5.** *Calls to `reduceByKey(⊕, prdd)` have deterministic outcomes if*

$$\text{lookup}(k, \text{reduceByKey}^D(\oplus, \text{part}(L))) = \text{reduce}(\oplus, \text{filterkey}(k, L))$$

for all list  $L$  of pairs, partitioning  $\text{part}$ , and key  $k$ .

**Proposition 4.** *Calls to `reduceByKey(⊕, prdd)` have deterministic outcomes iff calls to `reduce(⊕, rdd)` have deterministic outcomes.*

#### 4.5 Discussion

Our conditions for deterministic outcomes are more general than it appears. In addition to scalar data, such as integers, they are also applicable to RDDs containing non-scalar data, such as lists or sets. In our extended set of case studies, we will prove deterministic outcomes from a distributed Spark program using non-scalar data (Appendix B.3).

Corollary 1 gives necessary and sufficient conditions for calls to `aggregate` to have deterministic outcomes. Instead of checking whether `aggregate` computes the same result on all possible partitionings on any list for given  $z$ ,  $seq$ , and  $comb$ , the corollary, instead, allows us to investigate properties for all elements of  $\text{img}(\text{foldl}(seq, z)) \times \text{img}(\text{foldl}(seq, z))$  and  $\alpha \times \text{img}(\text{foldl}(seq, z))$ . Our precise conditions reduce the need of checking all partitionings to checking all elements of Cartesian products. It appears that deterministic outcomes from calls to combinators can be verified automatically. The problem, however, remains difficult for the following reasons:

- (a) The domain  $img(\text{foldl}(seq, z))$  can be infinite and in general not computable.
- (b) Even if  $\alpha$  and  $img(\text{foldl}(seq, z))$  are computable,  $seq$  and  $\oplus$  may not be computable. Naïvely enumerating elements in  $\alpha$  and  $img(\text{foldl}(seq, z))$  would not work.
- (c) Testing equality between elements of  $img(\text{foldl}(seq, z))$  can be undecidable.

Given  $seq :: \beta \times \alpha \rightarrow \beta$ , recall that  $img(\text{foldl}(seq, z))$  is a subset of  $\beta$ . A sound but incomplete way to avoid (a) in practice is to test the properties of  $\oplus$  on all elements of  $\beta$  instead. If a counterexample is found for some elements of  $\beta$ , the counterexample may not be valid in a real aggregate call because it may not belong to  $img(\text{foldl}(seq, z))$ . In practical cases, the sets  $\alpha$  and  $\beta$  are finite (such as machine integers) and equality between their elements is decidable. Even for such cases, checking if outcomes of aggregate are deterministic is still difficult since  $seq$  and  $\oplus$  might not terminate for some input. In many real Spark programs, however,  $seq$  and  $\oplus$  are very simple and thus computable (for instance, with only bounded loops or recursion). A semi-procedure to test these conditions might work on such practical examples.

## 5 Case Studies

We evaluated advantages of our PURESPARK specification on several case studies. In this section, we first analyze a Spark implementation of linear classification. Using the `treeAggregate` specification and its criteria for deterministic outcomes, we construct inputs yielding non-deterministic outcomes from the Spark implementation. Second, we analyze an implementation of a standard scaler and find a non-deterministic behavior there, too. Yet more case studies are provided in App. B.

### 5.1 Linear Classification

Linear classification is a well-known machine learning technique to classify data sets. Fix a set of *features*. A *data point* is a vector of numerical feature values. A *labeled data point* is a data point with a discrete label. Given a labeled data set, the *classification problem* is to classify (new) unlabeled data points by the labeled data set. A particularly useful subproblem is the *binary classification problem*. Consider, for instance, a data set of vital signs of some population; each data point is labeled by the diagnosis of a disease (positive or negative). The binary classification problem can be used to predict whether a person has the particular disease. Linear classification solves the binary classification problem by finding an optimal hyperplane to divide the labeled data points. After a hyperplane is obtained, linear classification predicts an unlabeled data point by the half-space containing the point. Logistic regression and linear Support Vector Machines (SVMs) are linear classification algorithms.

Consider a data set  $\{(\vec{x}_i, y_i) : 1 \leq i \leq n\}$  of data points  $\vec{x}_i \in \mathbb{R}^d$  labeled by  $y_i \in \{0, 1\}$ . Linear classification can be expressed as a numerical optimization problem:

$$\min_{\vec{w} \in \mathbb{R}^d} f(\vec{w}) \quad \text{with} \quad f(\vec{w}) = \xi R(\vec{w}) + \frac{1}{n} \sum_{i=1}^n L(\vec{w}; \vec{x}_i, y_i)$$

where  $\xi \geq 0$  is a *regularization parameter*,  $R(\vec{w})$  is a *regularizer*, and  $L(\vec{w}; \vec{x}_i, y_i)$  is a *loss function*. A vector  $\vec{w}$  corresponds to a hyperplane in the data point space. The vector  $\vec{w}_{opt}$  attaining the optimum hence classifies unlabeled data points with criteria defined by the objective function  $f(\vec{w})$ . Logistic regression and linear SVM are but

two instances of the optimization problem with objective functions defined by different regularizers and loss functions.

In the Spark machine learning library, the numerical optimization problem is solved by gradient descent. Very roughly, gradient descent finds a local minimum of  $f(\vec{w})$  by “walking” in the opposite direction of the gradient of  $f(\vec{w})$ . The mean of subgradients at data points is needed to compute the gradient of  $f(\vec{w})$ . The Spark machine learning library invokes `treeAggregate` to compute the mean. Floating-point addition is used as the `comb` parameter of the aggregate combinator. Since floating-point addition is not associative, we expect to observe non-deterministic outcomes (Proposition 1).

Consider the following three labeled data points:  $-10^{20}$  labeled with 1, 600 labeled with 0, and  $10^{20}$  labeled with 1. We create a 20-partition RDD with an equal number of the three labeled data points. The Spark machine learning library function `LogisticRegressionWithSGD.train` is used to generate a logistic regression model to predict the data points  $-10^{20}$ , 600, and  $10^{20}$  in each run. Among 49 runs, 19 of them classify the three data points into two different classes: the two positive data points are always classified in the same class, while the negative data point is in the other. The other 30 runs, however, classify all three data points into the same class. We observe similar predictions from `SVMWithSGD.train` with the same labeled data points. 37 out of 46 runs classify the data points into two different classes; the other 9 runs classify them into one class. Interestingly, the data points are always classified into two different classes by both logistic regression and linear SVM when the input RDD has only three partitions. As we expected from our analysis of the function, non-deterministic outcomes were witnessed in our Spark distributed environment.

## 5.2 Standard Scaler

Standardization of data sets is a common pre-processing step in machine learning. Many machine learning algorithms tend to perform better when the training set is similar to the standard normal distribution. In the Spark machine learning library, the class `StandardScaler` is provided to standardize data sets. The function `StandardScaler.fit` takes an RDD of raw data and returns an instance of `StandardScalerModel` to transform data points. Two transformations are available in `StandardScalerModel`. One standardizes a data point by mean, and the other normalizes by variance of raw data. If data points in raw data are transformed by mean, the transformed data points have the mean equal to 0. Similarly, if they are transformed by variance, the transformed data points have the variance 1.

The `StandardScaler` implementation uses `treeAggregate` to compute statistical information. It uses floating-point addition to combine means of raw data in different partitions. As in the previous use case, since floating-point addition is not associative, `StandardScaler` does not produce deterministic outcomes (Section 4.3). In our experiment, we create a 100-partition RDD with values  $-10^{20}$ , 600,  $10^{20}$  of the same number of occurrences. The mean of the data set is  $(-10^{20} \times n + 600 \times n + 10^{20} \times n) / (3n) = 200$  where  $n$  is the number of occurrences of each value. The data point 200 should therefore be after standardization transformed to 0. In 50 runs on the same data set in our distributed Spark platform, `StandardScaler` transforms 200 to a range of values from  $-944$  to 1142, validating our prediction of a non-deterministic outcome.

## 6 Related Work

**MapReduce modeling and optimization.** In the MapReduce (MR) computation, various cost and performance models have been proposed [26,17,15,32]. These models estimate the execution time and resource requirements of MR jobs. Karloff et al. developed a formal computation model for MR [20] and showed how a variety of algorithms can exploit the combination of sequential and parallel computation in MR. We are not aware of a similar work in the context of Spark. To the best of our knowledge, our work is the first to address the problem of formal, functional specification of Spark aggregation. Verifying the correctness of a MR program involves checking the commutativity and associativity of the reduce function. Xu et al. propose various semantic criteria to model commonly held assumptions on MR programs [29], including determinism, partition isolation, commutativity, and associativity of map/reduce combinators. Their empirical survey shows that these criteria are often overlooked by programmers and violated in practice. A recent survey [28] has found that a large number of industrial MR programs are, in fact, non-commutative. Recent work has proposed techniques for checking commutativity of bounded reducers automatically [12]. Because it is non-trivial to implement high-level algorithms using the MR framework, various approaches to compute optimized MR implementations have been proposed [16,23,25]. Emoto et al. [16] formalize the algebraic conditions using semiring homomorphism, under which an efficient program based on the generate-test-aggregate programming model can be specified in the MR framework. Given a monolithic *reduce* function, the work in [23] tries to decompose *reduce* into partial aggregation functions (similar to *seq* and *comb* in this paper) using program inversion techniques. MOLD [25] translates imperative Java code into MR code by transforming imperative loops into *fold* combinators using semantic-preserving program rewrite rules.

**Numerical Stability under MapReduce.** Several works try to scale up machine learning algorithms for large datasets using MapReduce [13,26]. To achieve numerically stable results across multiple runs [5,27], for example, preventing overflow, underflow and round-off errors due to finite-precision arithmetic, a variety of techniques are proposed [27]: generalizing sequential numerical stability techniques to distributed settings, shifting data values by constants, divide-and-conquer, etc. We showed that simulating machine learning algorithms using our specification enables early detection of points of numerical instability.

**Relational Query Optimization.** Relational query optimization is an extensively researched topic [11,19]: the goal is to obtain equivalent but more efficient query expressions by exploiting the algebraic properties of the constituent operators, for instance, join, select, together with statistics on relations and indices. For example, while inner joins commute independent of data, left joins commute only in specific cases. Query optimization for partitioned tables has received less attention [18,2]: because the key relational operators are not partition-aware, most work has focused on necessary but not sufficient conditions for query equivalence. In contrast, we investigate determinism of Spark aggregate expressions, constructed using partition-aware *seq* and *comb* combinators. We describe necessary and sufficient conditions under which these computations yield deterministic results independent of the data partitions.

**Deterministic Parallel Programming.** In order to enable deterministic-by-default parallel programming [7,10,8,9,21], researchers have developed several programming abstractions and logical specification languages to ensure that programs produce the same output for the same input independent of thread scheduling. For example, Deterministic Parallel Java [7,8] ensures exclusive writes to shared memory regions by means of verified, user-provided annotations over memory regions. In contrast, deterministic outcomes from Spark aggregation depend on algebraic properties like commutativity and associativity of *seq* and *comb* functions and their interplay

## 7 Conclusion

In this paper, we give a Haskell specification for various Spark aggregate combinators. We focus on aggregation of RDDs representing general sets, sets of pairs, and graphs. Based on our specification, we derive necessary and sufficient conditions that guarantee deterministic outcomes of the considered Spark aggregate combinators. We investigate several case studies and use the conditions to predict non-deterministic outcomes. Our executable specification can be used by developers for more detailed analysis and efficient development of distributed Spark programs. We also believe that our specifications are valuable resources for research communities to understand Spark better.

There are several future directions. The conditions for deterministic outcomes of aggregate combinators could be used for: (i) creating fully mechanized proofs for properties about data-parallel programs; (ii) developing automatic techniques for detecting non-deterministic outcomes of data-parallel programs; and (iii) synthesizing deterministic concurrent programs from sequential specifications. We have formalized the proofs of some crucial lemmas in Agda [4]. Using Scalaz [3], verified Haskell specifications can be translated to Spark programs to ensure determinism by construction.

## References

1. Apache Spark, <https://github.com/apache/spark>
2. IBM DB2 Version 9.7. Partitioned Tables, <https://ibm.biz/BdHyYR>
3. The Scalaz project, <https://github.com/scalaz>
4. PURESPARK, <https://github.com/guluchen/purespark>
5. Bennett, J., Grout, R., Pebay, P., Roe, D., Thompson, D.: Numerically stable, single-pass, parallel statistics algorithms. In: CLUSTER. pp. 1–8 (2009)
6. Bird, R.S.: An introduction to the theory of lists. In: the NATO Advanced Study Institute on Logic of programming and calculi of discrete design. pp. 5–42. Springer (1987)
7. Bocchino, Jr., R.L., Adve, V.S., Dig, D., Adve, S.V., Heumann, S., Komuravelli, R., Overbey, J., Simmons, P., Sung, H., Vakilian, M.: A type and effect system for deterministic parallel Java. In: OOPSLA. pp. 97–116 (2009)
8. Bocchino, Jr., R.L., Heumann, S., Honarmand, N., Adve, S.V., Adve, V.S., Welc, A., Shpeisman, T.: Safe nondeterminism in a deterministic-by-default parallel language. SIGPLAN Not. 46(1), 535–548 (2011)
9. Budimlic, Z., Burke, M.G., Cavé, V., Knobe, K., Lowney, G., Newton, R., Palsberg, J., Peixotto, D.M., Sarkar, V., Schlimbach, F., Tasirlar, S.: Concurrent collections. Scientific Programming 18(3-4), 203–217 (2010)
10. Burnim, J., Sen, K.: Asserting and checking determinism for multithreaded programs. Commun. ACM 53(6), 97–105 (2010)
11. Chaudhuri, S.: An overview of query optimization in relational systems. PODS '98 (1998)

12. Chen, Y., Hong, C., Sinha, N., Wang, B.: Commutativity of reducers. In: Proc. of TACAS'15. pp. 131–146. LNCS, Springer (2015)
13. Chu, C., Kim, S.K., Lin, Y., Yu, Y., Bradski, G.R., Ng, A.Y., Olukotun, K.: Map-Reduce for machine learning on multicore. In: NIPS. pp. 281–288 (2006)
14. Dean, J., Ghemawat, S.: MapReduce: A flexible data processing tool. *Commun. ACM* 53(1), 72–77 (2010)
15. Dörre, J., Apel, S., Lengauer, C.: Modeling and optimizing MapReduce programs. *Concurrency and Computation: Practice and Experience* 27(7), 1734–1766 (2015)
16. Emoto, K., Fischer, S., Hu, Z.: Generate, test, and aggregate: A calculation-based framework for systematic parallel programming with MapReduce. In: ESOP. pp. 254–273 (2012)
17. Herodotou, H., Babu, S.: Profiling, what-if analysis, and cost-based optimization of MapReduce programs. *Proceedings of the VLDB Endowment* 4(11), 1111–1122 (2011)
18. Herodotou, H., Borisov, N., Babu, S.: Query optimization techniques for partitioned tables. pp. 49–60. SIGMOD '11
19. Ioannidis, Y.E.: Query optimization. *ACM Comput. Surv.* 28(1), 121–123 (1996)
20. Karloff, H., Suri, S., Vassilvitskii, S.: A model of computation for MapReduce. In: SODA. pp. 938–948 (2010)
21. Leijen, D., Föhndrich, M., Burckhardt, S.: Prettier concurrency: Purely functional concurrent revisions. In: Haskell. pp. 83–94 (2011)
22. Leith, D., Clifford, P.: Convergence of distributed learning algorithms for optimal wireless channel allocation. In: IEEE Conference on Decision and Control. pp. 2980–2985 (2006)
23. Liu, C., Zhang, J., Zhou, H., McDirmid, S., Guo, Z., Moscibroda, T.: Automating distributed partial aggregation. In: SoCC. pp. 1:1–1:12 (2014)
24. Malewicz, G., Austern, M.H., Bik, A.J., Dehnert, J.C., Horn, I., Leiser, N., Czajkowski, G.: Pregel: A system for large-scale graph processing. In: ACM SIGMOD. pp. 135–146 (2010)
25. Radoi, C., Fink, S.J., Rabbah, R.M., Sridharan, M.: Translating imperative code to MapReduce. In: OOPSLA. pp. 909–927 (2014)
26. Sakr, S., Liu, A., Fayoumi, A.G.: The family of MapReduce and large-scale data processing systems. *ACM Comput. Surv.* 46(1), 11:1–11:44 (2013)
27. Tian, Y., Tatikonda, S., Reinwald, B.: Scalable and numerically stable descriptive statistics in SystemML. In: ICDE. pp. 1351–1359 (2012)
28. Xiao, T., Zhang, J., Zhou, H., Guo, Z., McDirmid, S., Lin, W., Chen, W., Zhou, L.: Non-terminism in MapReduce considered harmful? an empirical study on non-commutative aggregators in MapReduce programs. In: Companion Proceedings of ICSE. pp. 44–53 (2014)
29. Xu, Z., Hirzel, M., Rothermel, G.: Semantic characterization of MapReduce workloads. In: IISWC. pp. 87–97 (2013)
30. Zaharia, M., Chowdhury, M., Das, T., Dave, A., Ma, J., McCauly, M., Franklin, M.J., Shenker, S., Stoica, I.: Resilient distributed datasets: A fault-tolerant abstraction for in-memory cluster computing. In: NSDI. pp. 15–28 (2012)
31. Zaharia, M., Xin, R.S., Wendell, P., Das, T., Armbrust, M., Dave, A., Meng, X., Rosen, J., Venkataraman, S., Franklin, M.J., Ghodsi, A., Gonzalez, J., Shenker, S., Stoica, I.: Apache Spark: A unified engine for big data processing. *Commun. ACM* 59(11), 56–65 (Oct 2016)
32. Zhang, Z., Cherkasova, L., Verma, A., Loo, B.T.: Performance modeling and optimization of deadline-driven Pig programs. *ACM Trans. Auton. Adapt. Syst.* 8(3), 14:1–14:28 (2013)

## A Graph RDDs

Using RDDs, Spark provides a framework to analyze graphs distributively. In the Spark GraphX library, each vertex in a graph is designated by a `VertexId`, and associated with a vertex attribute. Each edge on the other hand is represented by `VertexIds` of its source and destination vertices. An edge is also associated with an edge attribute.

**type** `VertexId = Int`

**type** `VertexRDD  $\alpha$  = PairRDD VertexId  $\alpha$`

**type** `EdgeRDD  $\beta$  = RDD (VertexId, VertexId,  $\beta$ )`

**data** `GraphRDD  $\alpha$   $\beta$  = Graph { vertexRdd :: VertexRDD  $\alpha$ , edgeRdd :: EdgeRDD  $\beta$  }`

Let `graphRdd` be a graph RDD. Its vertex RDD (`vertexRdd graphRdd`) contains pairs of vertex identifiers and attributes. Different from conventional pair RDDs, each vertex identifier can appear at most once in the vertex RDD since a vertex is associated with exactly one attribute. If, for instance, two pairs with the same vertex identifier are generated during computation, their associated attributes must be merged to obtain a valid vertex RDD. The edge RDD (`edgeRdd graphRdd`) consists of triples of source and destination vertex identifiers, and edge attributes. Multi-edged directed graphs are allowed. In a graph RDD, the vertex and edge RDDs need to be consistent. That is, the source and destination vertex identifiers of any edge from the edge RDD must appear in the vertex RDD of the graph RDD.

The Spark GraphX library provides aggregate combinators for graph RDDs. We begin with an informal description of a slightly more general `aggregateMessagesWithActiveSet` combinator (Algorithm 1). The combinator takes functions `sendMsg` and `mergeMsg`, and a list `active` of vertices as its parameters. The list `active` determines *active* edges, that is, edges with source or destination vertex identifiers in `active`. For each active edge, the function `aggregateMessagesWithActiveSet` invokes `sendMsg` to send messages to its vertices. Messages sent to each vertex are merged by `mergeMsg`. Since a vertex is associated with at most one message after merging, the result is a valid vertex RDD.

```
foreach active edge e do
  | call sendMsg on e to send messages to vertices of e;
end
foreach vertex v receiving messages do
  | call mergeMsg to merge all messages sent to v;
end
return a vertex RDD with merged messages;
```

### Algorithm 1: `aggregateMessagesWithActiveSet`

Formally, the function `sendMsg` accepts source and destination vertex identifiers, attributes of the vertices, and the edge attribute of an edge as inputs. It sends messages to the source or destination vertex, both, or none. In our specification, `lookup` is used to obtain vertex attributes from a vertex RDD. We generate a pair RDD of vertex identifiers and messages by invoking `sendMsg` on every active edge. The messages associated with the same vertex are then merged by applying `reduceByKey` on the pair RDD. The resultant vertex RDD contains merged messages as vertex attributes. We call it a *message* RDD for clarity. Note that if a vertex from the input graph RDD does not receive any message, it is not present in the output message RDD. The combinator `aggregateMessages` in the Spark GraphX library is defined by `aggregateMessagesWithActiveSet`. It

invokes `aggregateMessagesWithActiveSet` by passing the list of all vertex identifiers as the active list. The combinator effectively applies `sendMsg` to every edge in a graph RDD.

```

aggregateMessagesWithActiveSet ::
  (VertexId → α → VertexId → α → β → [(VertexId, γ)])
  → (γ → γ → γ) → [VertexId] → GraphRDD α β → VertexRDD γ
aggregateMessagesWithActiveSet sendMsg mergeMsg active graphRdd =
  let isActive (srcId, dstId, _) = srcId 'elem' active || dstId 'elem' active
      vAttrs = concat (vertexRdd graphRdd)
      f edge = if isActive edge then
        let (srcId, dstId, edgeAttr) = edge
            srcAttr = fromJust (lookup srcId vAttrs)
            dstAttr = fromJust (lookup dstId vAttrs)
        in sendMsg srcId srcAttr dstId dstAttr edgeAttr
        else []
      pairRdd = map (concatMap f) (edgeRdd graphRdd)
  in reduceByKey mergeMsg pairRdd

aggregateMessages :: (VertexId → α → VertexId → α → β → [(VertexId, γ)])
  → (γ → γ → γ) → GraphRDD α β → VertexRDD γ
aggregateMessages sendMsg mergeMsg graphRdd =
  let vertices = concatMap (map fst) (vertexRdd graphRdd)
  in aggregateMessagesWithActiveSet sendMsg mergeMsg vertices graphRdd

```

Many graph algorithms perform fixed point computation. The Spark GraphX library hence provides a Pregel-like function to apply `aggregateMessages` on a graph RDD repetitively [24]. The Spark `pregel` function takes four input parameters `initMsg`, `vprog`, `sendMsg`, and `mergeMsg` (Algorithm 2). At initialization, it updates vertex attributes of the graph RDD by invoking `vprog` with the initial message `initMsg`. The `pregel` function then calls `aggregateMessages` to obtain a message RDD. If a vertex receives a message, its attribute is updated by `vprog` with the message. After updating vertex attributes, `pregel` obtains a new message RDD by invoking `aggregateMessagesWithActiveSet` with the active list equal to message-receiving vertices. Subsequently, only edges connecting to such vertices can send new messages.

```

foreach vertex v in G do
  | call vprog on v with initMsg to obtain its initial vertex attribute;
end
msgRdd ← call aggregateMessages on G;
while msgRdd is not empty do
  | foreach vertex v with message m in msgRdd do
    | call vprog on v with m to update its vertex attribute on G;
    end
    msgRdd ← call aggregateMessagesWithActiveSet with active equal to the vertices
    in msgRdd;
  end
return G;

```

**Algorithm 2:** `pregel`

We use several auxiliary functions to specify the Spark pregel function. Given a function computing an attribute from a vertex identifier and an attribute, the auxiliary function `mapVertexRDD` applies the function to every vertex in a vertex RDD and obtains another vertex RDD with new attributes. The `mapVertexRDD` function is used in `mapVertices` to update vertex attributes in graph RDDs. Moreover, recall that `aggregateMessagesWithActiveSet` returns a message RDD. The auxiliary function `joinGraph` updates a graph RDD with messages in a message RDD. For each vertex in the graph RDD, its attribute is joined with the message in the message RDD. If there is no message, the vertex attribute is left unchanged. The `pregel` function sets up the initial graph RDD by `mapVertices`. It then computes the initial message RDD by `aggregateMessages`. In each iteration, a new graph RDD is obtained by joining the graph RDD with a message RDD. `aggregateMessagesWithActiveSet` is then invoked to compute a new message RDD for the next iteration. The `pregel` function terminates when no more message is sent.

```
mapVertexRDD :: (VertexId → α → β) → VertexRDD α → VertexRDD β
mapVertexRDD f vRdd = map (map (λ(i, attr) → (i, f i attr))) vRdd
```

```
mapVertices :: (VertexId → α → γ) → GraphRDD α β → GraphRDD γ β
mapVertices updater gRdd = Graph {
  vertexRdd = mapVertexRDD updater (vertexRdd gRdd),
  edgeRdd = edgeRdd gRdd }
```

```
joinGraph :: (VertexId → α → γ → α) → GraphRDD α β
           → VertexRDD γ → GraphRDD α β
joinGraph joiner gRdd msgRdd = let assoc = concat msgRdd
  updt i attr = case lookup i assoc of Just v → joiner i attr v
  Nothing → attr
in mapVertices updt gRdd
```

```
pregel :: γ → (VertexId → α → γ → α) →
        (VertexId → α → VertexId → α → β → [(VertexId, γ)])
        → (γ → γ → γ) → GraphRDD α β → GraphRDD α β
pregel initMsg vprog sendMsg mergeMsg graphRdd =
  let initG = let init_f i attr = vprog i attr initMsg
  in mapVertices init_f graphRdd
  initMsgRdd = aggregateMessages sendMsg mergeMsg initG
  loop curG [] = curG
  loop curG msgRdd = let newG = joinGraph vprog curG msgRdd
  active = concatMap (map fst) msgRdd
  msgRdd' = aggregateMessagesWithActiveSet
  sendMsg mergeMsg active newG
  in loop newG msgRdd'
```

## A.1 Deterministic Aggregation in Graph Rdds

In this section, we explore necessary and sufficient conditions for aggregation in graph RDDs. In particular, we investigate deterministic outcomes of calls to the function

$\text{aggregateMessages}(send, \oplus, graphRdd)$  for  $send :: VertexID \times \alpha \times VertexID \times \alpha \times \beta \rightarrow [(VertexID, \gamma)]$ ,  $\oplus :: \gamma \times \gamma \rightarrow \gamma$ , and  $graphRdd :: GraphRDD \alpha \beta$ . We define deterministic outcomes first.

**Definition 6.** Calls to the function  $\text{aggregateMessages}(send, \oplus, graphRdd)$  have deterministic outcomes if for any two graph RDD representations of the same graph

$$graphRdd_1, graphRdd_2 :: GraphRDD \alpha \beta,$$

we have for all vertex identifiers  $v :: VertexID$ ,

$$\begin{aligned} \text{lookup}(v, \text{aggregateMessages}(send, \oplus, graphRdd_1)) = \\ \text{lookup}(v, \text{aggregateMessages}(send, \oplus, graphRdd_2)). \end{aligned}$$

The following proposition gives a sufficient condition for  $\text{aggregateMessages}$  to have deterministic outcomes.

**Proposition 5.** It holds that if calls to the function  $\text{reduceByKey}(\oplus, rdd)$  have deterministic outcomes, then calls to the function  $\text{aggregateMessages}(send, \oplus, graphRdd)$  also have deterministic outcomes.

## B Extended Set of Case Studies

This section of the appendix gives yet more case studies that we explored when analyzing Spark’s machine learning and graph libraries.

### B.1 Vertex Coloring

Let  $\Gamma = \{1, \dots, k\}$  denote the set of  $k$  colors. Given an undirected graph  $G = (V, E)$ , a  $k$ -coloring of  $G$  is a map  $C : V \rightarrow \Gamma$  such that  $C(v) \neq C(u)$  for any  $\{v, u\} \in E$ . In this case study, we will implement the Communication-Free Learning (CFL) algorithm [22] to find a  $k$ -coloring using the Spark GraphX library. Let  $0 < \beta < 1$ . The algorithm computes a  $k$ -coloring by iterations. We say a vertex  $v$  is *inactive* if all vertices adjacent to  $v$  have colors different from the color of  $v$ . Otherwise,  $v$  is *active*. At the  $n$ -th iteration, the CFL algorithm randomly chooses a color  $C_n(v) \in \Gamma$  by the color distribution  $P_n(v, \bullet)$  of  $v$ . The color distribution  $P_n(v, \bullet)$  is defined as follows. For  $n = 0$ ,  $P_0(v, c) = 1/k$  for all  $v \in V$  and  $c \in \Gamma$ . Each vertex hence chooses one of the  $k$  colors uniformly at random. For  $n > 0$ , let  $c = C_{n-1}(v)$  be the color of  $v$  in the previous iteration.

- If  $v$  is inactive, define  $P_n(v, c) = 1$  and  $P_n(v, d) = 0$  for  $d \neq c$ . Thus  $v$  does not change its color.
- Otherwise, define

$$P_n(v, d) = \begin{cases} (1 - \beta) \cdot P_{n-1}(v, c) & \text{if } d = c \\ (1 - \beta) \cdot P_{n-1}(v, d) + \beta/(k - 1) & \text{if } d \neq c \end{cases}$$

Thus  $v$  is more likely to choose a color different from  $c$ .

Observe that  $C_n$  stabilizes if and only if it is a  $k$ -coloring.

We implement the CFL algorithm using `pregel` in PURESPARK. For each vertex  $v$ , its attribute consists of the vertex color  $C_n(v)$ , the color distribution  $P_n(v, \bullet)$ , the vertex state (active or not), and a random number generator. As in Section B.3, an edge  $(u, v, \_)$  with  $u \geq v$  in an edge RDD represents  $\{u, v\} \in E$ . Given a graph RDD `graphRdd`, we construct its base graph `baseG` with initial vertex attributes.

```
initDist = map (\_ → 1.0 / fromIntegral k) [1..k]
```

```
baseG = mapVertices (\i _ → let (c, g) = randomR (1, k) (mkStdGen i)
                             in (c, initDist, True, g)) graphRdd
```

where `initDist` is the uniform distribution over  $k$  colors.

Consider the following `sendMsg` function:

```
sendMsg srcId (srcColor, _, srcActive, _) dstId (dstColor, _, dstActive, _) =
  if srcColor == dstColor then [(srcId, True), (dstId, True)]
  else (if srcActive then [(srcId, False)] else []) ++
       (if dstActive then [(dstId, False)] else [])
mergeMsg msg1 msg2 = msg1 || msg2
```

If the source and destination vertices of an edge have the same color, `sendMsg` sends `True` to both vertices to update vertex attributes. If they have different colors and the source vertex is active, `False` is sent to the source vertex. Similarly, `False` is sent to the destination vertex if the vertex is active. `mergeMsg` is the disjunction of messages. After applying `aggregateMessagesWithActiveSet` with `sendMsg` and `mergeMsg`, a vertex may receive a Boolean message. If a vertex receives `True`, it becomes active since one of its neighbors has the same color. Otherwise, the vertex becomes inactive.

We use `vprog` to update vertex attributes. For each vertex receiving a message, its vertex state, color, and color distribution are updated according to the CFL algorithm. The auxiliary function `sampleColor` chooses a color randomly by the color distribution. The helper function in `vprog` computes the color distribution  $P_n(v, \bullet)$  for the next iteration.

```
sampleColor dist p = let f (color, mass) weight =
                    (if m < p then succ color else color, m)
                    where m = mass + weight
                    in fst (foldl f (1, 0.0) dist)
```

```
vprog _ (c, dist, _, g) active = let helper (i, res) weight =
                                let decay = weight * (1 - beta)
                                    d = decay + (if c == i then 0 else beta / fromIntegral (numColors-1))
                                    e = if c == i then 1.0 else 0.0
                                in (succ i, if active then res ++ [d] else res ++ [e])
                                dist' = snd (foldl helper (1, []) dist)
                                (p, g') = random g
                                c' = if active then sampleColor dist' p else c
                                in (c', dist', active, g')
```

Finally, we invoke `pregel` to compute a  $k$ -coloring:

```
coloring = pregel True vprog sendMsg mergeMsg baseG
```

We test our executable Haskell specification on a typical Linux server. Since our Spark specification PURESPARK is faithful to Spark APIs, we realize it in the GraphX library with little manual effort. Our implementation works as intended on the distributed Spark platform.

## B.2 Connected Components

The Spark GraphX library implements a connected component algorithm for directed graphs. The documentation however does not explain what connected components are in directed graphs. We will find out what the implementation does here. Consider the following PURESPARK specification extracted from the Spark implementation:

```
connectedComponent graphRdd =
  let baseG = mapVertices ( $\lambda i \_ \rightarrow i$ ) graphRdd
      initMsg = maxBound :: Int
      sendMsg src srcA dst dstA _ =
        if srcA < dstA then [(dst, srcA)]
        else if dstA < srcA then [(src, dstA)]
        else []
      vprog _ attr msg = min attr msg
  in pregel initMsg vprog sendMsg min baseG
```

Given a graph RDD `graphRdd`, its base graph `baseG` is obtained by setting the attribute of a vertex to the identifier of the vertex. `sendMsg` compares the attributes of the source and destination vertices of an edge. The smaller attribute is sent to the vertex with the larger attribute. If both attributes are equal, no message is sent. If a number of messages are sent to a vertex, only the minimal message remains after applying `aggregateMessagesWithActiveSet` with `sendMsg` and `min`. When a vertex receives a message, its attribute is set to the minimum of its attribute and the message.

Consider a graph  $G = (V, E)$  with  $E \subseteq V \times V$ . We use  $attr(v)$  for the attribute of the vertex  $v \in V$ . Two vertices  $u$  and  $v$  are *linked* if  $(u, v) \in E$  or  $(v, u) \in E$ . Using our specification of `pregel`, it is not hard to see that the PURESPARK specification implements Algorithm 3. Note that the two for-each loops essentially propagate minimal attributes to linked vertices. When the set *active* is empty, the attributes of every linked vertices are equal and the algorithm terminates. We say two vertices  $u$  and  $v$  are *connected* if there are  $w_0 = u, w_1, \dots, w_k = v$  such that  $w_i$  and  $w_{i+1}$  are linked for  $0 \leq i < k$ . When `connectedComponent` terminates, connected vertices have the same attribute equal to the minimal vertex identifier among them. Hence the Spark implementation returns a graph RDD whose vertex attributes are the minimal vertex identifiers of connected vertices.

One can informally reason that the PURESPARK connected component specification has deterministic outcomes. Note that  $(VertexId, \min)$  is a commutative semigroup. This allows us to derive a similar proposition for `aggregateMessagesWithActiveSet`. The calls to `aggregateMessages` and `aggregateMessagesWithActiveSet` in `pregel` therefore have deterministic outcomes (Proposition 5). Examining the `vprog` in our connected component specification, the functions `mapVertices` and `joinGraph` also have deterministic outcomes. All potential sources of non-determinism in `pregel` have deterministic

```

attr(v) ← the vertex identifier of v;
active ← V;
while active ≠ ∅ do
  active' ← ∅;
  foreach v ∈ active do
    if attr(u) < attr(v) for some u linked with v then
      | send attr(u) to v and add v to active'
    if attr(v) < attr(u) for some u linked with v then
      | send attr(v) to u and add u to active'
    end
  foreach v ∈ active' do
    | attr(v) ← the minimal attribute sent to v
  end
  active ← active';
end

```

**Algorithm 3:** Connected Components

outcomes. The connected component specification consequently has deterministic outcomes. Experiments in a distributed Spark environment confirm our reasoning.

### B.3 Triangle Count

Let  $G = (V, E)$  be an undirected graph without self-loops or multiple edges. For  $u, v \in V$ ,  $\{u, v\} \in E$  denotes that  $u$  and  $v$  are adjacent. A *triangle* in  $G$  is formed by  $u, v, w \in V$  such that  $\{u, v\}, \{u, w\}, \{v, w\} \in E$ . Counting the number of triangles is important to, for example, network analysis. The Spark GraphX library implements the triangle counting algorithm using `aggregateMessages`.

In the GraphX implementation, an undirected graph is represented by a graph RDD where the source vertex identifier of every edge is greater than its destination vertex identifier. An edge  $\{u, v\} \in E$  with  $u > v$  is thus represented by  $(u, v, \_)$  in an edge RDD. Below is the PURESPARK specification extracted from the Spark GraphX implementation.

```

sendMsg src _ dst _ = [(dst, singleton src), (src, singleton dst)]
adjacentVRdd = aggregateMessages sendMsg (union) graphRdd

```

```

newGRdd = let adjacents = concat adjacentVRdd
          updt v _ = case lookup v adjacents of
                    Just adj → delete v adj
                    Nothing → empty
          in mapVertices updt graphRdd

```

```

sendMsg2 src srcA dst dstA _ =
  let num = size (intersection srcA dstA)
  in [(dst, num), (src, num)]
sumTriangles = aggregateMessages sendMsg2 (+) newGRdd

```

```

triangleCount = mapVertexRDD (λ_ y → quot y 2) sumTriangles

```

For each edge  $\{u, v\} \in E$ , `sendMsg` sends  $\{u\}$  and  $\{v\}$  to vertices  $v$  and  $u$  respectively. Multiple messages to a vertex are merged by union. After applying `aggregateMessages` with `sendMsg` and union, `adjacentVRdd` is a vertex RDD where the attribute of the vertex  $v$  is  $\{u : \{u, v\} \in E\}$ .

The implementation updates vertex attributes of the input graph to obtain `newGRdd`. If the set  $A$  of vertices adjacent to  $v$  is not empty, the attribute of  $v$  is updated to  $A \setminus \{v\}$ . If  $v$  does not have any adjacent vertices, its attribute is set to the empty set. Hence the attribute of a vertex in `newGRdd` contains its adjacent vertices but not itself. Recall that we assume the input graph does not have self-loops. A vertex cannot be adjacent to itself. Removing a vertex from the set of its adjacent vertices is redundant.

For each edge  $\{u, v\} \in E$  in `newGRdd`, `sendMsg2` sends the message  $|U \cap V|$  to  $u$  and  $v$  where  $U$  and  $V$  are the sets of vertices adjacent to  $u$  and  $v$  respectively. Observe that for every  $w \in U \cap V$ , we have  $\{w, u\}, \{w, v\}, \{u, v\} \in E$ . Let  $\Delta_{\{u, v\}}$  denote the number of triangles containing the edge  $\{u, v\}$ .  $\Delta_{\{u, v\}}$  is sent to both  $u$  and  $v$ . Messages are moreover merged by summation. Hence the attribute of each vertex  $v$  in `sumTriangles` is  $\sum_{\{u, v\} \in E} \Delta_{\{u, v\}}$ .

Now consider a vertex  $v$  in a triangle of  $u, v, w$ . The triangle is counted in both  $\Delta_{\{u, v\}}$  and  $\Delta_{\{w, v\}}$ . Since a triangle is always counted twice, the attribute given as  $\frac{1}{2} \sum_{\{u, v\} \in E} \Delta_{\{u, v\}}$  of vertex  $v$  in `triangleCount` is the number of triangles containing  $v$ . Both calls to `aggregateMessages` have deterministic outcomes because the algebras  $(\text{Set}, (\text{union}))$  and  $(\text{Int}, (+))$  are commutative semigroups (Propositions 4, 5, and Corollary 2).

## B.4 In-Degrees

The Spark GraphX library implements several graph algorithms using aggregation. We show how our specification helps to understand and analyze Spark programs utilizing aggregate combinators.

Let  $G = (V, E)$  with  $E \subseteq V \times V$  be a directed graph. We define the *in-degree* of a vertex  $v \in V$  as  $|\{(u, v) : (u, v) \in E\}|$ . The GraphX library uses the function `aggregateMessages` to compute in-degrees of vertices in a graph RDD. Consider the following Purespark specification for the GraphX implementation:

```
inDegrees graphRdd =
  let sendMsg _ _ dst _ = [(dst, 1)]
  in aggregateMessages sendMsg (+) graphRdd
```

By our specification, `aggregateMessages` invokes `sendMsg` on every edge in `graphRdd`. The `sendMsg` function sends the message 1 to the destination vertex of an edge. If several messages are sent to a vertex, they are summed up. Hence `inDegree` returns a vertex RDD where each vertex has the number of its incoming edges as the attribute. They are in-degrees of vertices in `graphRdd`. The call to `aggregateMessages` has a deterministic outcome because  $(\text{Int}, (+))$  is a commutative semigroup (Propositions 4, 5, and Corollary 2).

## C Missing Proofs

We start with proving the following auxiliary lemma.

**Lemma 4.**

$$\text{foldl}(f, z, p_1 \uplus p_2) = \text{foldl}(f, \text{foldl}(f, z, p_1), p_2) \quad (7)$$

*Proof.* By induction on the length of  $p_1$ .

– for  $p_1 = []$ :

$$\begin{aligned} \text{foldl}(f, \text{foldl}(f, z, []), p_2) &= \text{foldl}(f, z, p_2) && \text{(def. of foldl)} \\ &= \text{foldl}(f, z, [] \uplus p_2) && \text{(def. of } \uplus \text{)} \end{aligned}$$

– suppose the lemma holds for all  $p_1$  of length  $n$ . Now consider the list  $x : p_1$ . It follows that

$$\begin{aligned} \text{foldl}(f, z, x : p_1 \uplus p_2) &= \text{foldl}(f, f(z, x), p_1 \uplus p_2) && \text{(def. of foldl)} \\ &= \text{foldl}(f, \text{foldl}(f, f(z, x), p_1), p_2) && \text{(IH)} \\ &= \text{foldl}(f, \text{foldl}(f, z, x : p_1), p_2) && \text{(def. of foldl)} \quad \square \end{aligned}$$

In the following we use the following function:

$\text{aggregateList part z seq comb xs} = \text{aggregate}^D z \text{ seq comb (part xs)}$

**Lemma 5.** *The following are necessary (though not sufficient) conditions for a call  $\text{aggregate}(z, \text{seq}, \oplus, \text{part}(L))$  to be deterministic:*

1.  $z$  is the identity of  $\oplus$  on  $\gamma = \text{img}(\text{foldl}(\text{seq}, z))$ ,
2.  $\oplus$  is closed on  $\gamma$ ,
3.  $\oplus$  is commutative on  $\gamma$ , and
4.  $\oplus$  is associative on  $\gamma$ .

*Proof.* 1. We assume that  $\text{aggregate}(z, \text{seq}, \oplus, \text{part}(L))$  is deterministic and show that  $z$  is both the left and the right identity of  $\oplus$  on  $\gamma$ . First, assume the following partitioning:  $\text{part}_1(L) = [L]$ . From the assumption that the aggregate is deterministic, it follows that

$$\begin{aligned} \langle L \rangle &= \text{aggregateList}(\text{part}_1, z, \text{seq}, \oplus, L) \\ &= \text{foldl}(\oplus, z, [\langle L \rangle]) && \text{(def. of aggregateList)} \\ &= \text{foldl}(\oplus, z \oplus \langle L \rangle, []) && \text{(def. of foldl)} \\ &= z \oplus \langle L \rangle && \text{(def. of foldl)} \end{aligned}$$

Therefore,  $z$  is the left identity of  $\oplus$  on  $\gamma$ .

Second, assume the following partitioning:  $\text{part}_2(L) = [L, []]$ . From the assumption that the aggregate is deterministic, it follows that

$$\begin{aligned} \langle L \rangle &= \text{aggregateList}(\text{part}_2, z, \text{seq}, \oplus, L) \\ &= \text{foldl}(\oplus, z, [\langle L \rangle, \langle [] \rangle]) && \text{(def. of aggregateList)} \\ &= \text{foldl}(\oplus, z, [\langle L \rangle, z]) && \text{(def. of } \langle \cdot \rangle \text{ and foldl)} \\ &= \text{foldl}(\oplus, z \oplus \langle L \rangle, [z]) && \text{(def. of foldl)} \\ &= \text{foldl}(\oplus, \langle L \rangle, [z]) && (z \text{ is the left id. of } \oplus) \\ &= \text{foldl}(\oplus, \langle L \rangle \oplus z, []) && \text{(def. of foldl)} \\ &= \langle L \rangle \oplus z && \text{(def. of foldl)} \end{aligned}$$

Therefore,  $z$  is also the right identity of  $\oplus$  on  $\gamma$ .

2. We assume that  $\text{aggregate}(z, \text{seq}, \oplus, \text{rdd}(L))$  is deterministic and show that  $\oplus$  is closed on  $\gamma$ . First, we assume that  $L = p_1 \uplus p_2$  and consider the following partitioning:  $\text{part}(p_1 \uplus p_2) = [p_1, p_2]$ . From the assumption that the aggregate is deterministic, it follows that

$$\begin{aligned}
\langle p_1 \uplus p_2 \rangle &= \text{aggregateList}(\text{part}, z, \text{seq}, \oplus, L) \\
&= \text{foldl}(\oplus, z, [\langle p_1 \rangle, \langle p_2 \rangle]) && \text{(def. of aggregateList)} \\
&= \text{foldl}(\oplus, z \oplus \langle p_1 \rangle, [\langle p_2 \rangle]) && \text{(def. of foldl)} \\
&= \text{foldl}(\oplus, \langle p_1 \rangle, [\langle p_2 \rangle]) && (z \text{ is the id. of } \oplus) \\
&= \text{foldl}(\oplus, \langle p_1 \rangle \oplus \langle p_2 \rangle, []) && \text{(def. of foldl)} \\
&= \langle p_1 \rangle \oplus \langle p_2 \rangle && \text{(def. of foldl)}
\end{aligned}$$

Therefore  $\oplus$  is closed on  $\gamma$ .

3. We assume that  $\text{aggregate}(z, \text{seq}, \oplus, \text{rdd}(L))$  is deterministic and show that  $\oplus$  is commutative on  $\gamma$ . First, we assume that  $L = p_1 \uplus p_2$  and consider the following two partitionings:  $\text{part}_1(p_1 \uplus p_2) = [p_1, p_2]$  and  $\text{part}_2(p_1 \uplus p_2) = [p_2, p_1]$ . From the assumption that the aggregate is deterministic, it follows that

$$\begin{aligned}
&\text{aggregateList}(\text{part}_1, z, \text{seq}, \oplus, L) = \text{aggregateList}(\text{part}_2, z, \text{seq}, \oplus, L) \\
\iff &\text{foldl}(\oplus, z, [\langle p_1 \rangle, \langle p_2 \rangle]) = \text{foldl}(\oplus, z, [\langle p_2 \rangle, \langle p_1 \rangle]) && \text{(def. of aggregateList)} \\
\iff &\text{foldl}(\oplus, z \oplus \langle p_1 \rangle, [\langle p_2 \rangle]) = \text{foldl}(\oplus, z \oplus \langle p_2 \rangle, [\langle p_1 \rangle]) && \text{(def. of foldl)} \\
\iff &\text{foldl}(\oplus, \langle p_1 \rangle, [\langle p_2 \rangle]) = \text{foldl}(\oplus, \langle p_2 \rangle, [\langle p_1 \rangle]) && (z \text{ is the id. of } \oplus) \\
\iff &\text{foldl}(\oplus, \langle p_1 \rangle \oplus \langle p_2 \rangle, []) = \text{foldl}(\oplus, \langle p_2 \rangle \oplus \langle p_1 \rangle, []) && \text{(def. of foldl)} \\
\iff &\langle p_1 \rangle \oplus \langle p_2 \rangle = \langle p_2 \rangle \oplus \langle p_1 \rangle && \text{(def. of foldl)}
\end{aligned}$$

Therefore,  $\oplus$  is commutative on  $\gamma$ .

4. We assume that  $\text{aggregate}(z, \text{seq}, \oplus, \text{rdd}(L))$  is deterministic and show that  $\oplus$  is associative on  $\gamma$ . First, we assume that  $L = p_1 \uplus p_2 \uplus p_3$  and consider the following two partitionings:  $\text{part}_1(p_1 \uplus p_2 \uplus p_3) = [p_1, p_2, p_3]$  and  $\text{part}_2(p_1 \uplus p_2 \uplus p_3) = [p_2, p_3, p_1]$ . From the assumption that the aggregate is deterministic, it follows that

$$\begin{aligned}
&\text{aggregateList}(\text{part}_1, z, \text{seq}, \oplus, L) = \text{aggregateList}(\text{part}_2, z, \text{seq}, \oplus, L) \\
\iff &\text{foldl}(\oplus, z, [\langle p_1 \rangle, \langle p_2 \rangle, \langle p_3 \rangle]) = \text{foldl}(\oplus, z, [\langle p_2 \rangle, \langle p_3 \rangle, \langle p_1 \rangle]) && \text{(def. of aggregateList)} \\
\iff &\text{foldl}(\oplus, z \oplus \langle p_1 \rangle, [\langle p_2 \rangle, \langle p_3 \rangle]) = \text{foldl}(\oplus, z \oplus \langle p_2 \rangle, [\langle p_3 \rangle, \langle p_1 \rangle]) && \text{(def. of foldl)} \\
\iff &\text{foldl}(\oplus, \langle p_1 \rangle, [\langle p_2 \rangle, \langle p_3 \rangle]) = \text{foldl}(\oplus, \langle p_2 \rangle, [\langle p_3 \rangle, \langle p_1 \rangle]) && (z \text{ is the id. of } \oplus) \\
\iff &\text{foldl}(\oplus, \langle p_1 \rangle \oplus \langle p_2 \rangle, [\langle p_3 \rangle]) = \text{foldl}(\oplus, \langle p_2 \rangle \oplus \langle p_3 \rangle, [\langle p_1 \rangle]) && \text{(def. of foldl)} \\
\iff &\text{foldl}(\oplus, (\langle p_1 \rangle \oplus \langle p_2 \rangle) \oplus \langle p_3 \rangle, []) = \text{foldl}(\oplus, (\langle p_2 \rangle \oplus \langle p_3 \rangle) \oplus \langle p_1 \rangle, []) && \text{(def. of foldl)} \\
\iff &(\langle p_1 \rangle \oplus \langle p_2 \rangle) \oplus \langle p_3 \rangle = (\langle p_2 \rangle \oplus \langle p_3 \rangle) \oplus \langle p_1 \rangle && \text{(def. of foldl)} \\
\iff &(\langle p_1 \rangle \oplus \langle p_2 \rangle) \oplus \langle p_3 \rangle = \langle p_1 \rangle \oplus (\langle p_2 \rangle \oplus \langle p_3 \rangle) && \text{(comm. of } \oplus)
\end{aligned}$$

Therefore,  $\oplus$  is associative on  $\gamma$ . □

**Lemma 6.** For all functions  $h : [A] \rightarrow B$ , the following are equivalent:

1.  $h$  is a list homomorphism to  $(B, \odot, \perp)$ ,
2.  $\forall xss \in [[A]] : \text{foldl}(\odot, \perp, \text{map}(h, xss)) = h(\text{concat}(xss))$ .

*Proof.* (1  $\Rightarrow$  2): By induction on the length of  $xss$ :

- for  $xss = []$ :

$$\begin{aligned}
\text{foldl}(\odot, \perp, \text{map}(h, [])) &= \text{foldl}(\odot, \perp, []) && \text{(def. of map)} \\
&= \perp && \text{(def. of foldl)} \\
&= h([]) && \text{(assumption)} \\
&= h(\text{concat}([])) && \text{(def. of concat)}
\end{aligned}$$

- Consider the following induction hypothesis for  $xss_n$  of the length  $n$ :

$$\text{IH} : \text{foldl}(\odot, \perp, \text{map}(h, xss_n)) = h(\text{concat}(xss_n)). \quad (8)$$

For  $xss_n \# [xs]$  we proceed as follows:

$$\begin{aligned}
\text{foldl}(\odot, \perp, \text{map}(h, xss_n \# [xs])) &= \text{foldl}(\odot, \perp, \text{map}(h, xss_n) \# \text{map}(h, [xs])) && \text{(def. of map)} \\
&= \text{foldl}(\odot, \text{foldl}(\odot, \perp, \text{map}(h, xss_n)), \text{map}(h, [xs])) && \text{(Lemma 4)} \\
&= \text{foldl}(\odot, \text{foldl}(\odot, \perp, \text{map}(h, xss_n)), [h(xs)]) && \text{(def. of map)} \\
&= \text{foldl}(\odot, h(\text{concat}(xss_n)), [h(xs)]) && \text{(IH)} \\
&= h(\text{concat}(xss_n)) \odot h(xs) && \text{(def. of foldl)} \\
&= h(\text{concat}(xss_n) \# xs) && \text{(assumption)} \\
&= h(\text{concat}(xss_n \# [xs])) && \text{(def. of concat)}
\end{aligned}$$

(2  $\Rightarrow$  1): We prove that the two properties of a list homomorphism hold:

- From  $\text{foldl}(\odot, \perp, \text{map}(h, [])) = h(\text{concat}([]))$  it follows that  $h([]) = \perp$ .
- To prove that  $h(xs \# ys) = h(xs) \odot h(ys)$ , first we consider the list  $xss = [xs]$ :

$$\begin{aligned}
&\text{foldl}(\odot, \perp, \text{map}(h, [xs])) = h(\text{concat}([xs])) \\
\iff &\text{foldl}(\odot, \perp, [h(xs)]) = h(xs) && \text{(def. of map, def. of concat)} \\
\iff &\text{foldl}(\odot, \perp \odot h(xs), [ ]) = h(xs) && \text{(def. of foldl)} \\
\iff &\perp \odot h(xs) = h(xs) && \text{(def. of foldl)} \tag{9}
\end{aligned}$$

Then we consider the list  $xss = [xs, ys]$ :

$$\begin{aligned}
&\text{foldl}(\odot, \perp, \text{map}(h, [xs, ys])) = h(\text{concat}([xs, ys])) \\
\iff &\text{foldl}(\odot, \perp, [h(xs), h(ys)]) = h(xs \# ys) && \text{(def. of map, def. of concat)} \\
\iff &\text{foldl}(\odot, \perp \odot h(xs), [h(ys)]) = h(xs \# ys) && \text{(def. of foldl)} \\
\iff &\text{foldl}(\odot, (\perp \odot h(xs)) \odot h(ys), [ ]) = h(xs \# ys) && \text{(def. of foldl)} \\
\iff &(\perp \odot h(xs)) \odot h(ys) = h(xs \# ys) && \text{(def. of foldl)} \\
\iff &h(xs) \odot h(ys) = h(xs \# ys) && \text{((9))} \quad \square
\end{aligned}$$

**Lemma 1.** *Calls to  $\text{aggregate}(z, seq, \oplus, rdd)$  have deterministic outcomes iff:*

1.  $(\text{img}(\text{foldl}(seq, z)), \oplus, z)$  is a commutative monoid, and
2. for all lists  $p_1, p_2 :: [\alpha]$ ,  $\langle p_1 \# p_2 \rangle = \langle p_1 \rangle \oplus \langle p_2 \rangle$ .

*Proof.*  $\Rightarrow$ : (a) Proving 1: Follows from Lemma 5.

(b) Proving 2: consider the list  $xs \# ys$  and its partitioning  $\text{part}(xs \# ys) = [xs, ys]$ .

$$\begin{aligned}
&\text{aggregateList}(\text{part}, z, seq, \oplus, xs \# ys) = \langle xs \# ys \rangle && \text{(def. of det. aggregate)} \\
\iff &\text{foldl}(\oplus, z, [\langle xs \rangle, \langle ys \rangle]) = \langle xs \# ys \rangle && \text{(def. of aggregateList)} \\
\iff &\text{foldl}(\oplus, z \oplus \langle xs \rangle, [\langle ys \rangle]) = \langle xs \# ys \rangle && \text{(def. of foldl)} \\
\iff &\text{foldl}(\oplus, \langle xs \rangle, [\langle ys \rangle]) = \langle xs \# ys \rangle && (z \text{ is the id. of } \oplus) \\
\iff &\text{foldl}(\oplus, \langle xs \rangle \oplus \langle ys \rangle, [ ]) = \langle xs \# ys \rangle && \text{(def. of foldl)} \\
\iff &\langle xs \rangle \oplus \langle ys \rangle = \langle xs \# ys \rangle && \text{(def. of foldl)}
\end{aligned}$$

$\Leftarrow$ : Consider an arbitrary partitioning  $\text{part}(L)$  of  $L$  and its permutation  $\text{perm}$  s.t.  $L = \text{concat}(\text{perm}(\text{part}(L)))$ . From the definition of  $\langle \cdot \rangle$ , it follows that  $\langle [ ] \rangle = \text{foldl}(seq, z, [ ]) = z$ , and, therefore,  $\langle \cdot \rangle$  is a list homomorphism to  $(\text{img}(\text{foldl}(seq, z)), \oplus, z)$ . From Lemma 6 it follows that

$$\begin{aligned}
&\text{foldl}(\oplus, z, \text{map}(\langle \cdot \rangle, \text{perm}(\text{part}(L)))) = \langle \text{concat}(\text{perm}(\text{part}(L))) \rangle \\
\iff &\text{foldl}(\oplus, z, \text{map}(\langle \cdot \rangle, \text{perm}(\text{part}(L)))) = \langle L \rangle && \text{(def. of perm and part)} \\
\iff &\text{aggregateList}(\text{perm} \circ \text{part}, z, seq, \oplus, L) = \langle L \rangle && \text{(def. of aggregateList)}
\end{aligned}$$

Because  $\oplus$  is associative and commutative, it follows that  $\text{aggregateList}(\text{perm}_x \circ \text{part}, z, seq, \oplus, L) = \langle L \rangle$  for any  $\text{perm}_x$ . Therefore,  $\text{aggregate}(z, seq, \oplus, rdd(L))$  is deterministic.  $\square$

**Lemma 2.** Let  $\oplus$  be associative on  $\gamma = \text{img}(\text{foldl}(\text{seq}, z))$  and  $z$  be the identity of  $\oplus$  on  $\gamma$ . The following are equivalent:

$$1. \text{ for all lists } p_1, p_2 :: [\alpha], \quad \langle p_1 ++ p_2 \rangle = \langle p_1 \rangle \oplus \langle p_2 \rangle, \quad (2)$$

$$2. \text{ for all elements } d :: \alpha \text{ and } e :: \gamma, \quad \text{seq}(e, d) = e \oplus \text{seq}(z, d). \quad (3)$$

*Proof.* 1  $\implies$  2: This is a special case. We pick  $p_1$  such that  $\langle p_1 \rangle = e$  and  $p_2 = [d]$ . When we substitute into (2), we get

$$\langle p_1 ++ [d] \rangle = e \oplus \langle [d] \rangle. \quad (10)$$

For the left-hand side, according to Lemma 4, it holds that

$$\langle p_1 ++ [d] \rangle = \text{foldl}(\text{seq}, z, p_1 ++ [d]) = \text{foldl}(\text{seq}, \text{foldl}(\text{seq}, z, p_1), [d]) = \text{foldl}(\text{seq}, \langle p_1 \rangle, [d]). \quad (11)$$

After substitution, we get  $\text{foldl}(\text{seq}, e, [d])$ , which is (from the definition of  $\text{foldl}$ ) equal to  $\text{seq}(e, d)$ . For the right-hand side of (10), we just notice that  $\langle [d] \rangle = \text{foldl}(\text{seq}, z, [d]) = \text{seq}(z, d)$ .

2  $\implies$  1: Set  $x = \text{foldl}(\text{seq}, z, p_1) = \langle p_1 \rangle$  and substitute into (2) to obtain a new target for proving:

$$\begin{aligned} & \langle p_1 ++ p_2 \rangle = \langle p_1 \rangle \oplus \langle p_2 \rangle \\ \iff & \text{foldl}(\text{seq}, z, p_1 ++ p_2) = \langle p_1 \rangle \oplus \langle p_2 \rangle && \text{(def. of } \langle \cdot \rangle \text{)} \\ \iff & \text{foldl}(\text{seq}, \text{foldl}(\text{seq}, z, p_1), p_2) = \langle p_1 \rangle \oplus \langle p_2 \rangle && \text{(Lemma 4)} \\ \iff & \text{foldl}(x, \text{seq}, p_2) = x \oplus \langle p_2 \rangle && \text{(subst. of } x \text{)} \end{aligned} \quad (12)$$

We prove (12) using induction on the length  $n$  of  $p_2$ .

$n = 0$ : for  $p_2 = []$ , we get to prove the following:

$$\text{foldl}(\text{seq}, x, []) = x \oplus \text{foldl}(\text{seq}, z, []). \quad (13)$$

From the definition of  $\text{foldl}$ , we get an equivalent formula

$$x = x \oplus z, \quad (14)$$

which is true due to  $z$  being the identity of  $\oplus$  on  $\gamma$ .

$n = i + 1$ : We assume (12) holds for  $p_2$  of length  $i$ , i.e.

$$\text{IH : } \text{foldl}(\text{seq}, x, p_i) = x \oplus \text{foldl}(\text{seq}, z, p_i) \quad (15)$$

and prove that, for any  $h \in \alpha$ ,

$$\text{foldl}(\text{seq}, x, p_i ++ [h]) = x \oplus \text{foldl}(\text{seq}, z, p_i ++ [h]). \quad (16)$$

We do it in the following way:

$$\begin{aligned} & \text{foldl}(\text{seq}, x, p_i ++ [h]) \\ &= \text{foldl}(\text{seq}, \text{foldl}(\text{seq}, x, p_i), [h]) && \text{(Lemma 4)} \\ &= \text{foldl}(\text{seq}, \text{seq}(\text{foldl}(\text{seq}, x, p_i), h), []) && \text{(def. of foldl)} \\ &= \text{seq}(\text{foldl}(\text{seq}, x, p_i), h) && \text{(def. of foldl)} \\ &= \text{foldl}(\text{seq}, x, p_i) \oplus \text{seq}(z, h) && \text{(appl. of (3))} \\ &= (x \oplus \text{foldl}(\text{seq}, z, p_i)) \oplus \text{seq}(z, h) && \text{(IH)} \\ &= x \oplus (\text{foldl}(z, \text{seq}, p_i) \oplus \text{seq}(z, h)) && \text{(assoc. of } \oplus \text{)} \\ &= x \oplus \text{seq}(\text{foldl}(\text{seq}, z, p_i), h) && \text{(appl. of (3))} \\ &= x \oplus \text{foldl}(\text{seq}, \text{seq}(\text{foldl}(\text{seq}, z, p_i), h), []) && \text{(def. of foldl)} \\ &= x \oplus \text{foldl}(\text{seq}, \text{foldl}(\text{seq}, z, p_i), [h]) && \text{(def. of foldl)} \\ &= x \oplus \text{foldl}(\text{seq}, z, p_i ++ [h]) && \text{(Lemma 4)} \quad \square \end{aligned}$$

**Lemma 7.**

$$\text{reduce}(f, xs) = \text{reduce}'(f, xs) \quad (17)$$

where

**reducel' f xs = fromJust (foldl f' Nothing xs)**  
**where f' x y = case x of**  
**Nothing → Just y**  
**Just x' → Just (f x' y)**

*Proof.* by induction on the length of  $xs$ :

1. for  $xs = []$ , both reduce and reducel are undefined.
2. for  $xs = [x]$ :

$$\text{reducel}(f, [x]) = \text{foldl}(f, x, []) = x$$

and

$$\begin{aligned} \text{reducel}'(f, [x]) &= \text{fromJust}(\text{foldl}(f', \text{Nothing}, [x])) && \text{(def. of reducel')} \\ &= \text{fromJust}(\text{foldl}(f', f'(\text{Nothing}, x), [])) && \text{(def. of foldl)} \\ &= \text{fromJust}(\text{foldl}(f', \text{Just}(x), [])) && \text{(def. of } f') \\ &= \text{fromJust}(\text{Just}(x)) && \text{(def. of foldl)} \\ &= x && \text{(def. of fromJust)} \end{aligned}$$

3. assume the following induction hypothesis:

$$\text{reducel}(f, x : xs) = \text{reducel}'(f', x : xs) = R \tag{18}$$

We now prove that the lemma holds for  $x : xs ++ [a]$ . First, we compute the result for  $\text{reducel}(f, x : xs ++ [a])$ :

$$\begin{aligned} \text{reducel}(f, x : xs ++ [a]) &= \text{foldl}(f, x, xs ++ [a]) && \text{(def. of reducel)} \\ &= \text{foldl}(f, \text{foldl}(f, x, xs), [a]) && \text{(Lemma 4)} \\ &= \text{foldl}(f, \text{reducel}(f, x : xs), [a]) && \text{(def. of reducel)} \\ &= \text{foldl}(f, R, [a]) && \text{(IH)} \\ &= \text{foldl}(f, f(R, a), []) && \text{(def. of foldl)} \\ &= f(R, a) && \text{(def. of foldl)} \end{aligned}$$

We proceed by computing the result for  $\text{reducel}'(f', x : xs ++ [a])$ :

$$\begin{aligned} \text{reducel}'(f', x : xs ++ [a]) &= \text{fromJust}(\text{foldl}(f', \text{Nothing}, x : xs ++ [a])) && \text{(def. of reducel')} \\ &= \text{fromJust}(\text{foldl}(f', \text{foldl}(f', \text{Nothing}, x : xs), [a])) && \text{(Lemma 4)} \\ &= \text{fromJust}(\text{foldl}(f', f'(\text{foldl}(f', \text{Nothing}, x : xs), a), [])) && \text{(def. of foldl)} \\ &= \text{fromJust}(f'(\text{foldl}(f', \text{Nothing}, x : xs), a)) && \text{(def. of foldl)} \\ &\quad \langle f' \text{ is applied at least once on } x : xs \implies \text{the result of the nested foldl cannot be Nothing} \rangle \\ &= \text{fromJust}(\text{Just}(f(\text{fromJust}(\text{foldl}(f', \text{Nothing}, x : xs)), a))) && \text{(def. of } f') \\ &= f(\text{fromJust}(\text{foldl}(f', \text{Nothing}, x : xs)), a) && \text{(def. of fromJust)} \\ &= f(\text{reducel}'(f', x : xs), a) && \text{(def. of reducel')} \\ &= f(R, a) && \text{(IH)} \quad \square \end{aligned}$$

**Lemma 3.** *Calls to  $\text{reduce}(\oplus, \text{rdd})$  have deterministic outcomes iff calls to  $\text{aggregate}(\text{Nothing}, \text{seq}', \oplus, \text{rdd})$  have deterministic outcomes, where  $\text{seq}'$  and  $\oplus'$  are as follows:*

$\text{seq}' x y = \text{case } x \text{ of}$   $(\oplus') x y = \text{case } (x, y) \text{ of } (\text{Nothing}, y') \rightarrow y'$   
**Nothing → Just y**  $(x', \text{Nothing}) \rightarrow x'$   
**Just x' → Just (x'  $\oplus$  y)**  $(\text{Just } x', \text{Just } y') \rightarrow \text{Just } (x' \oplus y')$

*Proof.* We show that given the following definition of the function  $\text{reduce}''$ ,

$\text{reduce}'' :: (\alpha \rightarrow \alpha \rightarrow \alpha) \rightarrow \text{RDD } \alpha \rightarrow \alpha$   
 $\text{reduce}'' (\oplus) \text{rdd} = \text{fromJust } (\text{aggregate } \text{Nothing } \text{seq}' (\oplus') \text{rdd})$ ,

it holds that  $\text{reduce}''(\oplus, rdd) = \text{reduce}^D(\oplus, rdd)$  for all  $\oplus$  and  $rdd$ . In case  $rdd$  is a partitioning of an empty list, the result of both  $\text{reduce}'$  and  $\text{reduce}''$  is undefined. For a non-empty list:

$$\begin{aligned}
& \text{reduce}''(\oplus', xs : xss) \\
&= \text{fromJust}(\text{aggregate}(\text{Nothing}, seq', \oplus', xs : xss)) && \text{(def. of reduce}'') \\
&= \text{fromJust}(\text{foldl}(\oplus', \text{Nothing}, \text{map}(\lambda ys . \text{foldl}(seq', \text{Nothing}, ys), xs : xss))) && \text{(def. of aggregate)} \\
&\quad \langle \text{from the assumption on partitionings, no element of } xs : xss \text{ is empty} \rangle \\
&= \text{fromJust}(\text{foldl}(\oplus', \text{Nothing}, \text{map}(\lambda ys . \text{Just}(\text{fromJust}(\text{foldl}(seq', \text{Nothing}, ys))), xs : xss))) && \text{(def. of fromJust)} \\
&= \text{fromJust}(\text{foldl}(\oplus', \text{Nothing}, \text{map}(\lambda ys . \text{Just}(\text{reducel}(\oplus, ys)), xs : xss))) && \text{(Lemma 7)} \\
&= \text{fromJust}(\text{foldl}(\oplus', \text{Nothing}, \text{Just}(\text{reducel}(\oplus, xs)) : \text{map}(\lambda ys . \text{Just}(\text{reducel}(\oplus, ys)), xss))) && \text{(def. of map)} \\
&= \text{fromJust}(\text{foldl}(\oplus', \text{Nothing} \oplus' \text{Just}(\text{reducel}(\oplus, xs)), \text{map}(\lambda ys . \text{Just}(\text{reducel}(\oplus, ys)), xss))) && \text{(def. of foldl)} \\
&= \text{fromJust}(\text{foldl}(\oplus', \text{Just}(\text{reducel}(\oplus, xs)), \text{map}(\lambda ys . \text{Just}(\text{reducel}(\oplus, ys)), xss))) && \text{(def. of } \oplus') \\
&= \text{fromJust}(\text{Just}(\text{foldl}(\oplus, \text{reducel}(\oplus, xs), \text{map}(\lambda ys . \text{reducel}(\oplus, ys), xss)))) && \text{(def. of } \oplus') \\
&= \text{foldl}(\oplus, \text{reducel}(\oplus, xs), \text{map}(\lambda ys . \text{reducel}(\oplus, ys), xss)) && \text{(def. of fromJust)} \\
&= \text{reducel}(\oplus, \text{reducel}(\oplus, xs) : \text{map}(\lambda ys . \text{reducel}(\oplus, ys), xss)) && \text{(def. of reducel)} \\
&= \text{reducel}(\oplus, \text{map}(\lambda ys . \text{reducel}(\oplus, ys), xs : xss)) && \text{(def. of map)} \\
&= \text{reduce}^D(\oplus, xs : xss) && \text{(def. of reduce}^D) \quad \square
\end{aligned}$$

**Corollary 2.** *Calls to  $\text{reduce}(\oplus, rdd)$  have deterministic outcomes iff  $(\alpha, \oplus)$  is a commutative semigroup.*

*Proof.* From Lemma 3, it follows that we can investigate the function  $\text{aggregate}(\text{Nothing}, seq', \oplus', rdd)$  instead of  $\text{reduce}(\oplus, rdd)$ . From Corollary 1, we obtain that  $\text{aggregate}(\text{Nothing}, seq', \oplus', rdd)$  has deterministic outcome iff the following two conditions hold:

1.  $(\text{img}(\text{foldl}(seq', \text{Nothing})), \oplus', \text{Nothing})$  is a commutative monoid,
2.  $\forall d \in \alpha, e \in \text{img}(\text{foldl}(seq', \text{Nothing})) : seq'(e, d) = e \oplus' seq'(\text{Nothing}, d)$ .

We start with investigating condition 2:

– For the case  $e = \text{Nothing}$ :

$$\begin{aligned}
& seq'(e, d) = e \oplus' seq'(\text{Nothing}, d) \\
\iff & seq'(\text{Nothing}, d) = \text{Nothing} \oplus' seq'(\text{Nothing}, d) && \text{(subst. of } e = \text{Nothing)} \\
\iff & \text{Just}(d) = \text{Nothing} \oplus' \text{Just}(d) && \text{(def. of } seq') \\
\iff & \text{Just}(d) = \text{Just}(d) && \text{(def. of } \oplus')
\end{aligned}$$

– For the case  $e = \text{Just}(x)$ :

$$\begin{aligned}
& seq'(e, d) = e \oplus' seq'(\text{Nothing}, d) \\
\iff & seq'(\text{Just}(x), d) = \text{Just}(x) \oplus' seq'(\text{Nothing}, d) && \text{(subst. of } e = \text{Just}(x)) \\
\iff & \text{Just}(x \oplus d) = \text{Just}(x) \oplus' \text{Just}(d) && \text{(def. of } seq') \\
\iff & \text{Just}(x \oplus d) = \text{Just}(x \oplus d) && \text{(def. of } \oplus')
\end{aligned}$$

We can observe that the condition is a tautology. Therefore, the condition 1 is a sufficient and necessary condition for a call to  $\text{aggregate}(\text{Nothing}, seq', \oplus', rdd)$  to have a deterministic outcome.

We proceed by investigating the conditions for  $(\text{img}(\text{foldl}(seq', \text{Nothing})), \oplus', \text{Nothing})$  to be a commutative monoid. First, we observe that for  $\oplus : \alpha \times \alpha \rightarrow \alpha$ , it holds that  $\text{img}(\text{foldl}(seq', \text{Nothing})) = \text{Maybe}(\alpha)$ .

– *Identity:* From the definition,  $\text{Nothing}$  is the identity of  $\oplus'$ .

– *Commutativity:* From the definition,  $\oplus'$  is commutative iff  $\oplus$  is commutative.

– *Associativity:* Consider elements  $a, b, c \in \text{Maybe}(\alpha)$ . We explore when  $(a \oplus' b) \oplus' c = a \oplus' (b \oplus' c)$ :

- If any member of  $\{a, b, c\}$  is  $\text{Nothing}$ , the condition holds because  $\text{Nothing}$  is the (left and right) identity of  $\oplus'$ .
- For  $a = \text{Just}(a')$ ,  $b = \text{Just}(b')$ , and  $c = \text{Just}(c')$ , it holds that:

$$\begin{aligned}
& (\text{Just}(a) \oplus' \text{Just}(b)) \oplus' \text{Just}(c) = \text{Just}(a) \oplus' (\text{Just}(b) \oplus' \text{Just}(c)) \\
\iff & \text{Just}(a \oplus b) \oplus' \text{Just}(c) = \text{Just}(a) \oplus' \text{Just}(b \oplus c) && \text{(def. of } \oplus') \\
\iff & \text{Just}((a \oplus b) \oplus c) = \text{Just}(a \oplus (b \oplus c)) && \text{(def. of } \oplus')
\end{aligned}$$

Therefore,  $\oplus'$  is associative iff  $\oplus$  is associative.

– *Closed*: It is easy to observe that  $\oplus'$  is closed on  $\text{Maybe}(\alpha)$ .

From the previous conditions, we infer that  $\text{aggregate}(\text{Nothing}, \text{seq}', \oplus', \text{rdd})$  has deterministic outcome iff  $(\alpha, \oplus)$  is a commutative semiring.  $\square$

**Proposition 1.** *Calls to  $\text{treeAggregate}(z, \text{seq}, \oplus, \text{rdd})$  have deterministic outcomes iff calls to  $\text{aggregate}(z, \text{seq}, \oplus, \text{rdd})$  have deterministic outcomes.*

*Proof.*  $\Rightarrow$ : Consider the following function:

```
dividel :: [\alpha] -> ([\alpha], \alpha, \alpha, [\alpha])
dividel x1:x2:xs = ([], x1, x2, xs) .
```

Obviously, `dividel` is one possible way how `dividel` can function. We further consider the following modification of `apply`:

```
applyl :: (\beta -> \beta -> \beta) -> [\beta] -> \beta
applyl comb [r] = r
applyl comb [r, r'] = comb r r'
applyl comb rs = let (ls', l', r', rs') = dividel rs in applyl comb (ls' ++ [comb l' r'] ++ rs')
```

After inlining `dividel` to `applyl`, we can modify it to obtain yet further modification:

```
applyl' :: (\beta -> \beta -> \beta) -> [\beta] -> \beta
applyl' comb [r] = r
-- applyl' comb [r, r'] = comb r r'
applyl' comb r1:r2:rs = applyl' comb ((comb r1 r2):rs)
```

Note that the case for a list of length 2 is redundant now. Clearly it holds that  $\text{applyl}'(f, xs) = \text{reducel}(f, xs)$ . If we substitute `reducel` for `apply` in the definition of `treeAggregate`, and further use the property of a partitioning that it is never an empty list, we obtain the definition of `aggregate`.

$\Leftarrow$ : From Lemma 5, it follows that  $\oplus$  is associative and commutative. Therefore, any sequence of `dividel`-`apply` operations in `apply` will yield the same outcome as if we consider the (deterministic) `dividel`.  $\square$

**Proposition 2.** *Calls to  $\text{treeReduce}(\oplus, \text{rdd})$  have deterministic outcomes iff calls to  $\text{reduce}(\oplus, \text{rdd})$  have deterministic outcomes.*

*Proof.* Follows the same structure as the proof of Proposition 1.  $\square$

When inferring conditions for a deterministic outcome of the call to `aggregateByKey`, we make use of the following auxiliary function:

```
aggregateWithKey :: \alpha -> \gamma -> (\gamma -> \beta -> \gamma) -> (\gamma -> \gamma -> \gamma) -> PairRDD \alpha \beta -> \gamma
aggregateWithKey k z seq comb pairRdd =
  let select p = key p == k
      vrdd = filter (not . null)
              (map ((map value) . (filter select)) pairRdd)
  in aggregate z seq comb vrdd
```

We also use the following version of `aggregateByKey` with the partitioning given explicitly:

```
aggregateListByKey :: [(\alpha, \beta)] -> [(\alpha, \beta)] -> \gamma -> (\gamma -> \beta -> \gamma)
                    -> (\gamma -> \gamma -> \gamma) -> [(\alpha, \beta)] -> PairRDD \alpha \gamma
aggregateListByKey part z mergeComb mergeValue list = aggregateByKey z mergeComb mergeValue (part list)
```

**Lemma 8.** *It holds that*

$$\text{lookUp}(k, \text{aggregateByKey}(z, \text{seq}, \oplus, \text{prdd})) = \text{aggregateWithKey}(k, z, \text{seq}, \oplus, \text{prdd}),$$

where `lookUp` searches the first value with a given key in an RDD:

$$\text{lookUp}(k, xss) = \text{head}_z(\text{concat}(\text{map}(\text{map}(\text{value} \circ \text{filterkey } k), xss))),$$

and  $\text{head}_z$  returns  $z$  when the input is empty.

*Proof.* To avoid too many parentheses, we use curried functions for the proof of this lemma. We need a number of additional lemmas. The following property allows one to swap filterkey  $k$  and foldl (mergeBy ( $\oplus$ )) []:

$$\text{filterkey } k \circ \text{foldl (mergeBy } (\oplus) \text{) []} = \text{foldl (mergeBy } (\oplus) \text{) []} \circ \text{filterkey } k. \quad (19)$$

The next property says that, given a key  $k$  and a binary operator ( $\odot$ ), filtering the list with  $k$  and performing foldl (mergeBy( $\odot$ )) [] gives you a single value:

$$\text{head}_z \circ \text{map value} \circ \text{foldl (mergeBy } (\odot) \text{) []} \circ \text{filterkey } k = \text{foldl } \odot \ z \circ \text{map value} \circ \text{filterkey } k, \quad (20)$$

where  $\text{head}_z$  returns  $z$  when the input is empty. Finally, in the equation below, given a RDD and any binary operator ( $\odot$ ), the LHS computes foldl (mergeBy ( $\odot$ )) [] on each partition, pick those with key  $k$ , and concatenates their values. The RHS filters the values with key  $k$ , and computes foldl ( $\odot$ )  $z$  for each partition.

$$\begin{aligned} & \text{concat} \circ \text{map (map value} \circ \text{filterkey } k \circ \text{foldl (mergeBy } (\odot) \text{) [])} \\ &= \text{map (foldl } (\odot) \ z) \circ \text{filter (not} \circ \text{null)} \circ \text{map (map value} \circ \text{filterkey } k). \end{aligned} \quad (21)$$

All the lemmas above can be proved by induction. The proof of this lemma follows:

$$\begin{aligned} & \text{lookUp } k \circ \text{aggregateByKey } z \ (\otimes) \ (\oplus) \\ &= \text{head}_z \circ \text{concat} \circ \text{map (map value} \circ \text{filterkey } k) \circ \text{repartition} \circ \\ & \quad \text{foldl (mergeBy } (\oplus) \text{) []} \circ \text{concat} \circ \text{map (foldl (mergeBy } (\otimes) \text{) [])} \circ \text{perm} \quad (\text{def. of aggregateByKey}) \\ &= \text{head}_z \circ \text{map value} \circ \text{filterkey } k \circ \text{foldl (mergeBy } (\oplus) \text{) []} \circ \\ & \quad \text{concat} \circ \text{map (foldl (mergeBy } (\otimes) \text{) [])} \circ \text{perm} \quad (\text{naturality}) \\ &= \text{head}_z \circ \text{map value} \circ \text{foldl (mergeBy } (\oplus) \text{) []} \circ \text{filterkey } k \circ \\ & \quad \text{concat} \circ \text{map (foldl (mergeBy } (\otimes) \text{) [])} \circ \text{perm} \quad (\text{by (19)}) \\ &= \text{foldl } (\oplus) \ z \circ \text{map value} \circ \text{filterkey } k \circ \text{concat} \circ \text{map (foldl (mergeBy } (\otimes) \text{) [])} \circ \text{perm} \quad (\text{by (20)}) \\ &= \text{foldl } (\oplus) \ z \circ \text{concat} \circ \text{map (map value} \circ \text{filterkey } k \circ \text{foldl (mergeBy } (\otimes) \text{) [])} \circ \text{perm} \quad (\text{naturality}) \\ &= \text{foldl } (\oplus) \ z \circ \text{map (foldl } (\otimes) \ z) \circ \text{filter (not} \circ \text{null)} \circ \text{map (map value} \circ \text{filterkey } k) \circ \text{perm} \quad (\text{by (21)}) \\ &= \text{foldl } (\oplus) \ z \circ \text{map (foldl } (\otimes) \ z) \circ \text{perm} \circ \text{filter (not} \circ \text{null)} \circ \text{map (map value} \circ \text{filterkey } k) \quad (\text{naturality}) \\ &= \text{aggregateWithKey } k \ z \ (\otimes) \ (\oplus) \quad (\text{def. of aggregateWithKey}) \quad \square \end{aligned}$$

**Proposition 3.** *Calls to aggregateByKey( $z, seq, \oplus, prdd$ ) have deterministic outcomes iff calls to aggregate( $z, seq, \oplus, rdd$ ) have deterministic outcomes.*

*Proof.* From Lemma 8, it follows that aggregateByKey( $z, seq, \oplus, prdd$ ) has deterministic outcome iff for all keys  $k \in \alpha$  and partitionings  $part$ :

$$\text{aggregateWithKey}(k, z, seq, \oplus, part(L)) = \text{foldl}(z, seq, \text{filterkey}(k, L)). \quad (22)$$

From the definition of aggregateWithKey, we infer that this is equivalent to

$$\begin{aligned} & \text{aggregate}(z, seq, \oplus, part(\text{filterkey}(k, L))) = \text{foldl}(z, seq, \text{filterkey}(k, L)) \\ \iff & \text{aggregate}(z, seq, \oplus, part(L')) = \text{foldl}(z, seq, L'), \quad (\text{subst. } L' = \text{filterkey}(k, L)) \end{aligned}$$

which is the condition for aggregate( $z, seq, \oplus, part(L')$ ) to have a deterministic outcome.  $\square$

Consider the following function.

```
reduceWithKey ::  $\alpha \rightarrow (\beta \rightarrow \beta \rightarrow \beta) \rightarrow \text{PairRDD } \alpha \ \beta \rightarrow \beta$ 
reduceWithKey k mergeValue pairRdd =
  let select p = key p == k
      vrdd = filter (not . null)
              (map ((map value) . (filter select)) pairRdd)
  in reduce mergeValue vrdd
```

**Lemma 9.** *It holds that*

$$\text{lookup}(k, \text{reduceByKey}(\oplus, prdd)) = \text{reduceWithKey}(k, \oplus, prdd).$$

*Proof.* Similar to that of Lemma 8. □

**Proposition 4.** *Calls to  $\text{reduceByKey}(\oplus, prdd)$  have deterministic outcomes iff calls to  $\text{reduce}(\oplus, rdd)$  have deterministic outcomes.*

*Proof.* Follows the same structure as the proof of Proposition 3. □

**Proposition 5.** *It holds that if calls to the function  $\text{reduceByKey}(\oplus, rdd)$  have deterministic outcomes, then calls to the function  $\text{aggregateMessages}(send, \oplus, graphRdd)$  also have deterministic outcomes.*

*Proof.* When  $\text{reduceByKey}$  has deterministic outcome, then it holds (from definition) that for all vertices  $v \in \text{VertexID}$ , lists  $L \in [\alpha]$ , and partitionings  $part$ :

$$\text{lookup}(v, \text{reduceListWithKey}(part, \oplus, L)) = \text{reducel}(\oplus, \text{filterkey}(v, L)).$$

When applying  $\text{lookup}(v, \text{aggregateMessages}(send, \oplus, graphRdd(V, E)))$ , the result will be the same as if the lookup is applied to the last line of function  $\text{aggregateMessagesWithActiveSet}$ :

**lookup**( $v, \text{reduceByKey}(\oplus, \text{pairRdd})$ ).

Since  $\text{reduceByKey}(\oplus, \text{pairRdd})$  has deterministic outcome, it follows that

$$\text{lookup}(v, \text{reduceByKey}(\oplus, \text{pairRdd})) = \text{reducel}(\oplus, \text{filterkey}(v, \text{pairRdd})). \tag{23}$$

This is a sufficient condition to conclude that  $\text{aggregateMessages}(send, \oplus, graphRdd(V, E))$  has a deterministic outcome. □