

An Executable Sequential Specification for SPARK Aggregation

Yu-Fang Chen¹, Chih-Duo Hong¹, **Ondřej Lengál**^{1,2},
Shin-Cheng Mu¹, Nishant Sinha³, Bow-Yaw Wang¹

¹Academia Sinica, Taiwan

²Brno University of Technology, Czech Republic

³IBM Research, India

19 May 2017 (NETYS'17)

Current Trends in CAV (Computer-Aided Verification)

Current Trends in CAV (Computer-Aided Verification):

- verification of HW

Current Trends in CAV (Computer-Aided Verification)

Current Trends in CAV (Computer-Aided Verification):

- verification of HW
- verification of sequential programs:
 - ▶ w/ integers
 - ▶ w/ floats
 - ▶ w/ heap manipulation
 - ▶ ...

Current Trends in CAV (Computer-Aided Verification)

Current Trends in CAV (Computer-Aided Verification):

- verification of HW
- verification of sequential programs:
 - ▶ w/ integers
 - ▶ w/ floats
 - ▶ w/ heap manipulation
 - ▶ ...
- verification of concurrent programs
 - ▶ mutual exclusion protocols
 - ▶ concurrent data structures
 - ▶ ...

Current Trends in Computer Science


Current Trends in Computer Science:

- machine learning, deep neural networks, IoT, smart-*, ...

Current Trends in Computer Science

Current Trends in Computer Science:


- machine learning, deep neural networks, IoT, smart-*, ...
- big data & cluster computing
 - ▶ sciences
 - ▶ AI
 - ▶ advertising analysis
 - ▶ data mining
 - ▶ biology
 - ▶ search engines
 - ▶ market value:
 - expected over \$50 billion by 2020



Current Trends in Computer Science

Current Trends in Computer Science:

- machine learning, deep neural networks, IoT, smart-*, ...
- big data & cluster computing
 - ▶ sciences
 - ▶ AI
 - ▶ advertising analysis
 - ▶ data mining
 - ▶ biology
 - ▶ search engines
 - ▶ market value:
 - expected over \$50 billion by 2020
- Can verification be applied here, too?



Programming for Big Data:

Programming for Big Data:

- many repeating tasks
 - ▶ distribution of data on nodes

Programming for Big Data:

- many repeating tasks
 - ▶ distribution of data on nodes
 - ▶ collection of computed results

Programming for Big Data:

- many repeating tasks
 - ▶ distribution of data on nodes
 - ▶ collection of computed results
- \rightsquigarrow frameworks that create abstraction over the communication

Programming for Big Data:

- many repeating tasks
 - ▶ distribution of data on nodes
 - ▶ collection of computed results
- \rightsquigarrow frameworks that create abstraction over the communication
- web services providing easy-to-setup computation using these (Amazon, Microsoft, IBM, ...)

Programming for Big Data:

- many repeating tasks
 - ▶ distribution of data on nodes
 - ▶ collection of computed results
- \rightsquigarrow frameworks that create abstraction over the communication
- web services providing easy-to-setup computation using these (Amazon, Microsoft, IBM, ...)
- examples:
 - ▶ Hadoop MAPREDUCE
 - ▶ PIG
 - ▶ HIVE
 - ▶ **Apache SPARK**

New verification problems:

New verification problems:

- verification of correctness of the frameworks

New verification problems:

- verification of correctness of the frameworks
- verification of correctness of user programs
 - ▶ **correctness**: checking special properties

Apache SPARK

Apache SPARK:

- *successor* of Hadoop MAPREDUCE
 - ▶ claims to be up to $100\times$ faster due to *in-memory* computation

Apache SPARK

Apache SPARK:

- *successor* of Hadoop MAPREDUCE
 - ▶ claims to be up to $100\times$ faster due to *in-memory* computation
- a relaxed *fault tolerant* model
 - ▶ sub-results are recomputed upon faults

Apache SPARK

Apache SPARK:

- *successor* of Hadoop MAPREDUCE
 - ▶ claims to be up to $100\times$ faster due to *in-memory* computation
- a relaxed *fault tolerant* model
 - ▶ sub-results are recomputed upon faults
- *lazy evaluation* semantics

Apache SPARK:

- *successor* of Hadoop MAPREDUCE
 - ▶ claims to be up to $100\times$ faster due to *in-memory* computation
- a relaxed *fault tolerant* model
 - ▶ sub-results are recomputed upon faults
- *lazy evaluation* semantics
- contains libraries for
 - ▶ processing graphs
 - ▶ streaming computation
 - ▶ machine learning
 - ▶ SQL-based database computation
 - ▶ ...

Apache SPARK


RDD—Resilient Distributed Dataset:

- the principal data abstraction

Apache SPARK

RDD—Resilient Distributed Dataset:

- the principal data abstraction



Computation in SPARK

Computation in SPARK

- map-style

- ▶ **map, filter**

```
RDD rdd = ...
```

```
RDD newRdd = rdd.map( $\lambda x . x * 2$ )
```

Computation in SPARK

Computation in SPARK

■ map-style

▶ **map, filter**

```
RDD rdd = ...  
RDD newRdd = rdd.map( $\lambda x . x * 2$ )
```

■ **aggregation**

▶ **aggregate, reduce**

▶ **treeAggregate, treeReduce**

```
RDD rdd = ...  
int sum = rdd.reduce( $\lambda x y . x + y$ )
```


Computation in SPARK

Computation in SPARK

■ map-style

▶ **map**, **filter**

```
RDD rdd = ...  
RDD newRdd = rdd.map( $\lambda x . x * 2$ )
```

■ **aggregation**

▶ **aggregate**, **reduce**

▶ **treeAggregate**, **treeReduce**

```
RDD rdd = ...  
int sum = rdd.reduce( $\lambda x y . x + y$ )
```

■ combined

▶ **aggregateByKey**, **reduceByKey**

```
PairRDD rdd = ...  
PairRDD sum = rdd.reduceByKey( $\lambda x y . x + y$ )
```

Formal specification of SPARK aggregation

Formal specification of SPARK aggregation

- SPARK is written in SCALA (bindings to Java, Python, ...)
 - ▶ multi-paradigm programming language
 - ▶ both imperative and functional code (side-effects)
 - ▶ unclear semantics

Formal specification of SPARK aggregation

Formal specification of SPARK aggregation

- SPARK is written in SCALA (bindings to Java, Python, . . .)
 - ▶ multi-paradigm programming language
 - ▶ both imperative and functional code (side-effects)
 - ▶ unclear semantics
- documentation is not clear about requirements of functions

Formal specification of SPARK aggregation

Formal specification of SPARK aggregation

- SPARK is written in SCALA (bindings to Java, Python, . . .)
 - ▶ multi-paradigm programming language
 - ▶ both imperative and functional code (side-effects)
 - ▶ unclear semantics
- documentation is not clear about requirements of functions
- our contribution:
 - ▶ **PURESPARK**: a specification of aggregation functions in HASKELL
 - purely functional language
 - executable specification
 - suitable for formal reasoning (e.g. AGDA)

Formal specification of SPARK aggregation

Formal specification of SPARK aggregation

- SPARK is written in SCALA (bindings to Java, Python, . . .)
 - ▶ multi-paradigm programming language
 - ▶ both imperative and functional code (side-effects)
 - ▶ unclear semantics
- documentation is not clear about requirements of functions
- our contribution:
 - ▶ **PURESPARK**: a specification of aggregation functions in HASKELL
 - purely functional language
 - executable specification
 - suitable for formal reasoning (e.g. AGDA)
 - ▶ correctness requirements on aggregation functions (next slide)

Formal specification of SPARK aggregation

Formal specification of SPARK aggregation

- SPARK is written in SCALA (bindings to Java, Python, . . .)
 - ▶ multi-paradigm programming language
 - ▶ both imperative and functional code (side-effects)
 - ▶ unclear semantics
- documentation is not clear about requirements of functions
- our contribution:
 - ▶ **PURESPARK**: a specification of aggregation functions in HASKELL
 - purely functional language
 - executable specification
 - suitable for formal reasoning (e.g. AGDA)
 - ▶ correctness requirements on aggregation functions (next slide)
 - ▶ analysis of case studies—finding numeric instability in ML library

Commutativity in SPARK aggregation

Commutativity in SPARK aggregation


`aggregate(seq, comb, \perp , rdd)`



d1	d2	d3	d4
----	----	----	----



d5	d6	d7	d8	d9	d10	d11
----	----	----	----	----	-----	-----







d12	d13	d14
-----	-----	-----

Commutativity in SPARK aggregation

Commutativity in SPARK aggregation

`aggregate(seq, comb, \perp , rdd)`



`foldl(seq, \perp , [d1, d2, d3, d4])`

$\rightsquigarrow r_A$

`foldl(seq, \perp , [d5, d6, d7, d8, d9, d10, d11])`

$\rightsquigarrow r_B$

`foldl(seq, \perp , [d12, d13, d14])`

$\rightsquigarrow r_C$

Commutativity in SPARK aggregation

Commu

$$\mathbf{fold1} :: (B \rightarrow A \rightarrow B) \rightarrow B \rightarrow [A] \rightarrow B$$
$$\mathbf{fold1}(\mathit{seq}, \perp, []) = \perp$$
$$\mathbf{fold1}(\mathit{seq}, \perp, x : xs) = \mathbf{fold1}(\mathit{seq}, \mathit{seq}(\perp, x), xs)$$

d1	d2	d3	d4
----	----	----	----

d5	d6	d7	d8	d9	d10	d11
----	----	----	----	----	-----	-----





d12	d13	d14
-----	-----	-----

$$\mathbf{fold1}(\mathit{seq}, \perp, [d_1, d_2, d_3, d_4]) \rightsquigarrow r_A$$
$$\mathbf{fold1}(\mathit{seq}, \perp, [d_5, d_6, d_7, d_8, d_9, d_{10}, d_{11}]) \rightsquigarrow r_B$$
$$\mathbf{fold1}(\mathit{seq}, \perp, [d_{12}, d_{13}, d_{14}]) \rightsquigarrow r_C$$

Commutativity in SPARK aggregation

Commutativity in SPARK aggregation

`aggregate(seq, comb, \perp , rdd)`



`foldl(seq, \perp , [d1, d2, d3, d4])`

$\rightsquigarrow r_A$

`foldl(seq, \perp , [d5, d6, d7, d8, d9, d10, d11])`

$\rightsquigarrow r_B$

`foldl(seq, \perp , [d12, d13, d14])`

$\rightsquigarrow r_C$






■ collected results:

rC	rA	rB
----	----	----

Commutativity in SPARK aggregation

Commutativity in SPARK aggregation

`aggregate(seq, comb, \perp , rdd)`



`foldl(seq, \perp , [d1, d2, d3, d4])`

$\rightsquigarrow r_A$

`foldl(seq, \perp , [d5, d6, d7, d8, d9, d10, d11])`

$\rightsquigarrow r_B$

`foldl(seq, \perp , [d12, d13, d14])`

$\rightsquigarrow r_C$

■ collected results:


rC	rA	rB
----	----	----

nondeterministic!

Commutativity in SPARK aggregation

Commutativity in SPARK aggregation

`aggregate(seq, comb, \perp , rdd)`



d1 d2 d3 d4

d5 d6 d7 d8 d9 d10 d11

d12 d13 d14

`foldl(seq, \perp , [d1, d2, d3, d4])`

$\rightsquigarrow r_A$

`foldl(seq, \perp , [d5, d6, d7, d8, d9, d10, d11])`

$\rightsquigarrow r_B$

`foldl(seq, \perp , [d12, d13, d14])`

$\rightsquigarrow r_C$

■ collected results: r_C r_A r_B

nondeterministic!

`foldl(comb, \perp , [rC, rA, rB])`

\rightsquigarrow result


Commutativity in SPARK aggregation

Two sources of nondeterminism:


Commutativity in SPARK aggregation

Two sources of nondeterminism:

1. Partitioning into RDD



a)




b)


Commutativity in SPARK aggregation

Two sources of nondeterminism:

1. Partitioning into RDD




a)



b)


2. Order in which nodes send partial results




Commutativity in SPARK aggregation

Two sources of nondeterminism:

1. Partitioning into RDD




a)



b)

2. Order in which nodes send partial results



aggregate can yield different results!!!

Commutativity in SPARK aggregation

Example of a nondeterministic aggregation

aggregate(*seq*, *comb*, \perp , *rdd*)

$$\text{seq}(\text{acc}, x) = \text{acc} + x$$

$$\text{comb}(\text{lhs}, \text{rhs}) = \text{rhs} + \text{rhs} \quad (\text{typo})$$

$$\perp = 0$$

Commutativity in SPARK aggregation


Example of a nondeterministic aggregation

`aggregate(seq, comb, \perp , rdd)`

$seq(acc, x) = acc + x$

$comb(lhs, rhs) = rhs + rhs$ (typo)

$\perp = 0$



Commutativity in SPARK aggregation


Example of a nondeterministic aggregation

aggregate(seq, comb, \perp , rdd)

$$\text{seq}(\text{acc}, x) = \text{acc} + x$$

$$\text{comb}(\text{lhs}, \text{rhs}) = \text{rhs} + \text{rhs} \quad (\text{typo})$$

$$\perp = 0$$



$$\text{foldl}(\text{seq}, 0, [d_1, d_2, d_3, d_4]) = (((0 + d_1) + d_2) + d_3) + d_4 \rightsquigarrow r_A$$

$$\text{foldl}(\text{seq}, 0, [d_5, d_6, d_7, d_8, d_9, d_{10}, d_{11}]) = \dots \rightsquigarrow r_B$$

$$\text{foldl}(\text{seq}, 0, [d_{12}, d_{13}, d_{14}]) = \dots \rightsquigarrow r_C$$

Commutativity in SPARK aggregation


Example of a nondeterministic aggregation

aggregate(seq, comb, \perp , rdd)

$$\text{seq}(\text{acc}, x) = \text{acc} + x$$

$$\text{comb}(\text{lhs}, \text{rhs}) = \text{rhs} + \text{rhs} \quad (\text{typo})$$

$$\perp = 0$$



$$\text{foldl}(\text{seq}, 0, [d_1, d_2, d_3, d_4]) = (((0 + d_1) + d_2) + d_3) + d_4 \rightsquigarrow r_A$$

$$\text{foldl}(\text{seq}, 0, [d_5, d_6, d_7, d_8, d_9, d_{10}, d_{11}]) = \dots \rightsquigarrow r_B$$

$$\text{foldl}(\text{seq}, 0, [d_{12}, d_{13}, d_{14}]) = \dots \rightsquigarrow r_C$$

Collecting results:

■ : $\text{foldl}(\text{comb}, 0, [r_C, r_A, r_B]) = 2r_B$

■ : $\text{foldl}(\text{comb}, 0, [r_A, r_B, r_C]) = 2r_C$

Commutativity in SPARK aggregation


Example of a nondeterministic aggregation

aggregate(seq, comb, \perp , rdd)

$$\text{seq}(\text{acc}, x) = \text{acc} + x$$

$$\text{comb}(\text{lhs}, \text{rhs}) = \text{rhs} + \text{rhs} \quad (\text{typo})$$

$$\perp = 0$$



$$\text{foldl}(\text{seq}, 0, [d_1, d_2, d_3, d_4]) = (((0 + d_1) + d_2) + d_3) + d_4 \rightsquigarrow r_A$$

$$\text{foldl}(\text{seq}, 0, [d_5, d_6, d_7, d_8, d_9, d_{10}, d_{11}]) = \dots \rightsquigarrow r_B$$

$$\text{foldl}(\text{seq}, 0, [d_{12}, d_{13}, d_{14}]) = \dots \rightsquigarrow r_C$$

Collecting results:

■ : $\text{foldl}(\text{comb}, 0, [r_C, r_A, r_B]) = 2r_B$

■ : $\text{foldl}(\text{comb}, 0, [r_A, r_B, r_C]) = 2r_C$

$$2r_B \neq 2r_C$$

Commutativity in SPARK aggregation

Commutativity of aggregate

Definition

A call

aggregate(*seq*, *comb*, \perp , *rdd*)

is **commutative** iff

aggregate(*seq*, *comb*, \perp , *rdd*(*L*)) = **foldl**(*seq*, \perp , *L*)

for every partitioning *rdd*(*L*) of *L*.

Commutativity in SPARK aggregation

Commutativity of aggregate

Definition

A call

aggregate(*seq*, *comb*, \perp , *rdd*)

is **commutative** iff

aggregate(*seq*, *comb*, \perp , *rdd*(*L*)) = **foldl**(*seq*, \perp , *L*)

for every partitioning *rdd*(*L*) of *L*.

- i.e., **aggregate** is an implementation of **foldl**

Commutativity in SPARK aggregation

Commutativity of aggregate

Definition

A call

aggregate(*seq*, *comb*, \perp , *rdd*)

is **commutative** iff

aggregate(*seq*, *comb*, \perp , *rdd*(*L*)) = **foldl**(*seq*, \perp , *L*)

for every partitioning *rdd*(*L*) of *L*.

- i.e., **aggregate** is an implementation of **foldl**
- if a call to **aggregate** is commutative:
 - ▶ the call is deterministic
 - ▶ when analyzing the program, we can assume one partitioning

Commutativity in SPARK aggregation

Conditions for commutative aggregate

Theorem

Consider rdd of elements of type \mathbb{T} and $\perp \in \mathbb{R}$. A call

aggregate(seq, comb, \perp , rdd)

is *commutative* iff

- 1 $(\text{img}(\text{foldl}(\text{seq}, \perp)), \text{comb}, \perp)$ is a *commutative monoid* and
- 2 for all $d \in \mathbb{T}$ and $e \in \text{img}(\text{foldl}(\text{seq}, \perp))$, it holds that
$$\text{seq}(e, d) = \text{comb}(e, \text{seq}(z, d)).$$

Commutativity in SPARK aggregation

Conditions for commutative aggregate

Theorem

Consider rdd of elements of type \mathbb{T} and $\perp \in \mathbb{R}$. A call

aggregate(seq, comb, \perp , rdd)

is *commutative* iff

- 1 $(\text{img}(\text{foldl}(seq, \perp)), \text{comb}, \perp)$ is a *commutative monoid* and
- 2 for all $d \in \mathbb{T}$ and $e \in \text{img}(\text{foldl}(seq, \perp))$, it holds that
$$seq(e, d) = \text{comb}(e, seq(z, d)).$$

Safe approximation:





... is *commutative* if

- 1 $(\mathbb{R}, \text{comb}, \perp)$ is a *commutative monoid* and
- 2 for all $d \in \mathbb{T}$ and $e \in \mathbb{R}$, it holds that
$$seq(e, d) = \text{comb}(e, seq(z, d)).$$

Commutativity in SPARK aggregation

SPARK reduce





`reduce(comb, rdd)`



Commutativity in SPARK aggregation

SPARK reduce

reduce(*comb*, *rdd*)



reduce1(*comb*, [d_1, d_2, d_3, d_4])

$\rightsquigarrow r_A$


reduce1(*comb*, [$d_5, d_6, d_7, d_8, d_9, d_{10}, d_{11}$])

$\rightsquigarrow r_B$

reduce1(*comb*, [d_{12}, d_{13}, d_{14}])

$\rightsquigarrow r_C$

$\text{reducel} :: (A \rightarrow A \rightarrow A) \rightarrow [A] \rightarrow A$
 $\text{reducel}(\text{comb}, x : xs) = \text{foldl}(\text{comb}, x, xs)$



$\text{reducel}(\text{comb}, [d_1, d_2, d_3, d_4]) \rightsquigarrow r_A$




$\text{reducel}(\text{comb}, [d_5, d_6, d_7, d_8, d_9, d_{10}, d_{11}]) \rightsquigarrow r_B$

$\text{reducel}(\text{comb}, [d_{12}, d_{13}, d_{14}]) \rightsquigarrow r_C$

Commutativity in SPARK aggregation

SPARK reduce

reduce(*comb*, *rdd*)



reducer1(*comb*, [d_1, d_2, d_3, d_4])

$\rightsquigarrow r_A$

reducer1(*comb*, [$d_5, d_6, d_7, d_8, d_9, d_{10}, d_{11}$])

$\rightsquigarrow r_B$

reducer1(*comb*, [d_{12}, d_{13}, d_{14}])

$\rightsquigarrow r_C$





■ collected results:



Commutativity in SPARK aggregation

SPARK reduce

reduce(*comb*, *rdd*)



reduce1(*comb*, [d_1, d_2, d_3, d_4])

$\rightsquigarrow r_A$

reduce1(*comb*, [$d_5, d_6, d_7, d_8, d_9, d_{10}, d_{11}$])

$\rightsquigarrow r_B$

reduce1(*comb*, [d_{12}, d_{13}, d_{14}])

$\rightsquigarrow r_C$

■ collected results:




nondeterministic!

Commutativity in SPARK aggregation

SPARK reduce

reduce(*comb*, *rdd*)



d1 d2 d3 d4

d5 d6 d7 d8 d9 d10 d11

d12 d13 d14

reduce1(*comb*, [d_1, d_2, d_3, d_4])

$\rightsquigarrow r_A$

reduce1(*comb*, [$d_5, d_6, d_7, d_8, d_9, d_{10}, d_{11}$])

$\rightsquigarrow r_B$

reduce1(*comb*, [d_{12}, d_{13}, d_{14}])

$\rightsquigarrow r_C$

■ collected results:

rC rA rB

nondeterministic!

reduce1(*comb*, [r_C, r_A, r_B])

\rightsquigarrow **result**

Commutativity in SPARK aggregation

Conditions for commutative `reduce`

`reduce(comb, rdd)`

Commutativity in SPARK aggregation

Conditions for commutative `reduce`

`reduce(comb, rdd)`

- via reduction to `aggregate` (using the `Maybe` monad):

`aggregate(seq2, comb2, Nothing, rdd)`

$seq_2(x, y) = \text{case } x \text{ of}$
Nothing \rightarrow **Just** y
Just $v \rightarrow$ **Just** $comb(v, y)$

$comb_2(x, y) = \text{case } (x, y) \text{ of}$
(Nothing, y') $\rightarrow y'$
(x', Nothing) $\rightarrow x'$
(Just v₁, Just v₂) \rightarrow **Just** $comb(v_1, v_2)$

Commutativity in SPARK aggregation

Conditions for commutative `reduce`

`reduce(comb, rdd)`

- via reduction to `aggregate` (using the `Maybe` monad):

`aggregate(seq2, comb2, Nothing, rdd)`

$seq_2(x, y) = \text{case } x \text{ of}$
 $\text{Nothing} \rightarrow \text{Just } y$
 $\text{Just } v \rightarrow \text{Just } comb(v, y)$

$comb_2(x, y) = \text{case } (x, y) \text{ of}$
 $(\text{Nothing}, y') \rightarrow y'$
 $(x', \text{Nothing}) \rightarrow x'$
 $(\text{Just } v_1, \text{Just } v_2) \rightarrow \text{Just } comb(v_1, v_2)$

Theorem

Consider `rdd` of elements of type \mathbb{T} . A call

`reduce(comb, rdd)`

is *commutative* iff $(\mathbb{T}, comb)$ is a *commutative semigroup*.

Commutativity in SPARK aggregation

SPARK `treeAggregate` and `treeReduce`

- first stage same as `aggregate` and `reduce`

Commutativity in SPARK aggregation

SPARK `treeAggregate` and `treeReduce`

- first stage same as **`aggregate`** and **`reduce`**
- second stage is performed concurrently in a **binary tree structure**

Commutativity in SPARK aggregation

SPARK `treeAggregate` and `treeReduce`

- first stage same as `aggregate` and `reduce`
- second stage is performed concurrently in a `binary tree structure`

Theorem

- *`treeAggregate(seq, comb, \perp , rdd)` is commutative iff `aggregate(seq, comb, \perp , rdd)` is commutative.*
- *`treeReduce(comb, rdd)` is commutative iff `reduce(comb, rdd)` is commutative.*

Commutativity in SPARK aggregation

SPARK `aggregateByKey` and `reduceByKey`

- work on PairRDDs: elements are (k, v)

Commutativity in SPARK aggregation

SPARK `aggregateByKey` and `reduceByKey`

- work on PairRDDs: elements are (k, v)
- produce (again) a PairRDD
 - ▶ for every k , at most one pair $(k, result)$
 - ▶ $result$ = the output of **aggregate** on elements with the key k

Commutativity in SPARK aggregation

SPARK `aggregateByKey` and `reduceByKey`

- work on PairRDDs: elements are (k, v)
- produce (again) a PairRDD
 - ▶ for every k , at most one pair $(k, result)$
 - ▶ $result$ = the output of **aggregate** on elements with the key k

Theorem

- ***aggregateByKey**(seq, comb, \perp , rdd) is commutative iff **aggregate**(seq, comb, \perp , rdd) is commutative.*
- ***reduceByKey**(comb, rdd) is commutative iff **reduce**(comb, rdd) is commutative.*

Discussion

Conditions for deterministic aggregation:

- 1 $(\text{img}(\mathbf{foldl}(\text{seq}, \perp)), \text{comb}, \perp)$ is a **commutative monoid** and
- 2 for all $d \in \mathbb{T}$ and $e \in \text{img}(\mathbf{foldl}(\text{seq}, \perp))$, it holds that
$$\text{seq}(e, d) = \text{comb}(e, \text{seq}(z, d)).$$

Discussion

Conditions for deterministic aggregation:

- 1 $(img(\mathbf{foldl}(seq, \perp)), comb, \perp)$ is a **commutative monoid** and
- 2 for all $d \in \mathbb{T}$ and $e \in img(\mathbf{foldl}(seq, \perp))$, it holds that
$$seq(e, d) = comb(e, seq(z, d)).$$

are general:

- apart from scalar data (e.g. integers), they also work for non-scalar (e.g. lists, sets)

Discussion

Conditions for deterministic aggregation:

- 1 $(img(\mathbf{foldl}(seq, \perp)), comb, \perp)$ is a **commutative monoid** and
- 2 for all $d \in \mathbb{T}$ and $e \in img(\mathbf{foldl}(seq, \perp))$, it holds that
$$seq(e, d) = comb(e, seq(z, d)).$$

are general:

- apart from scalar data (e.g. integers), they also work for non-scalar (e.g. lists, sets)

issues:

- $img(\mathbf{foldl}(seq, \perp))$ can be infinite, in general not computable

Discussion

Conditions for deterministic aggregation:

- 1 $(img(\mathbf{foldl}(seq, \perp)), comb, \perp)$ is a **commutative monoid** and
- 2 for all $d \in \mathbb{T}$ and $e \in img(\mathbf{foldl}(seq, \perp))$, it holds that
$$seq(e, d) = comb(e, seq(z, d)).$$

are general:

- apart from scalar data (e.g. integers), they also work for non-scalar (e.g. lists, sets)

issues:

- $img(\mathbf{foldl}(seq, \perp))$ can be infinite, in general not computable
- seq and $comb$ can be general functions \rightsquigarrow may not terminate

Discussion

Conditions for deterministic aggregation:

- 1 $(img(\mathbf{foldl}(seq, \perp)), comb, \perp)$ is a **commutative monoid** and
- 2 for all $d \in \mathbb{T}$ and $e \in img(\mathbf{foldl}(seq, \perp))$, it holds that
$$seq(e, d) = comb(e, seq(z, d)).$$

are general:

- apart from scalar data (e.g. integers), they also work for non-scalar (e.g. lists, sets)

issues:

- $img(\mathbf{foldl}(seq, \perp))$ can be infinite, in general not computable
- seq and $comb$ can be general functions \rightsquigarrow may not terminate
- testing the universal equality may be undecidable

Discussion

Conditions for deterministic aggregation:

- 1 $(img(\mathbf{foldl}(seq, \perp)), comb, \perp)$ is a **commutative monoid** and
- 2 for all $d \in \mathbb{T}$ and $e \in img(\mathbf{foldl}(seq, \perp))$, it holds that
$$seq(e, d) = comb(e, seq(z, d)).$$

are general:

- apart from scalar data (e.g. integers), they also work for non-scalar (e.g. lists, sets)

issues:

- $img(\mathbf{foldl}(seq, \perp))$ can be infinite, in general not computable
- seq and $comb$ can be general functions \rightsquigarrow may not terminate
- testing the universal equality may be undecidable
- result for MAPREDUCE [Chen, Hong, Sinha, Wang; TACAS'15]
 - ▶ $\mathbb{N}, +, \times$, control(loop-free): undecidable

Commutativity in SPARK aggregation

Case studies:

- manual evaluation of SPARK ML library

Commutativity in SPARK aggregation

Case studies:

- manual evaluation of SPARK ML library
- many functions use `floats`

Commutativity in SPARK aggregation

Case studies:

- manual evaluation of SPARK ML library
- many functions use `floats`
- found a redundancy in the SPARK Graph library

Commutativity in SPARK aggregation

Our contribution:

- **PURESPARK**: a specification of aggregation functions in HASKELL
 - ▶ purely functional language
 - ▶ executable specification
 - ▶ suitable for formal reasoning (e.g. AGDA)

Commutativity in SPARK aggregation

Our contribution:

- **PURESPARK**: a specification of aggregation functions in HASKELL
 - ▶ purely functional language
 - ▶ executable specification
 - ▶ suitable for formal reasoning (e.g. AGDA)
- correctness requirements on aggregation functions

Commutativity in SPARK aggregation

Our contribution:

- **PURESPARK**: a specification of aggregation functions in HASKELL
 - ▶ purely functional language
 - ▶ executable specification
 - ▶ suitable for formal reasoning (e.g. AGDA)
- correctness requirements on aggregation functions
- analysis of case studies—finding numeric instability and redundancy in ML library

Commutativity in SPARK aggregation

Our contribution:

- **PURESPARK**: a specification of aggregation functions in HASKELL
 - ▶ purely functional language
 - ▶ executable specification
 - ▶ suitable for formal reasoning (e.g. AGDA)
- correctness requirements on aggregation functions
- analysis of case studies—finding numeric instability and redundancy in ML library
- also extended to **aggregate** over graphs