

Translating imperative code to MapReduce

Cosmin Rădoi ¹ Stephen J. Fink ²
Rodric Rabbah ² Manu Sridharan ³

¹University of Illinois

²IBM T.J. Watson Research Center

³Samsung Research America

background: What is MapReduce?

Jeffrey Dean and Sanjay Ghemawat.

MapReduce: simplified data processing on large clusters. OSDI'04

background: What is MapReduce?

simple programming model for processing big data on a cluster

Jeffrey Dean and Sanjay Ghemawat.

MapReduce: simplified data processing on large clusters. OSDI'04

background: What is MapReduce?

simple programming model for processing big data on a cluster

advantages:

- ▶ fault-tolerance
- ▶ elastic scaling
- ▶ integration with distributed file systems

Jeffrey Dean and Sanjay Ghemawat.

MapReduce: simplified data processing on large clusters. OSDI'04

background: What is MapReduce?

simple programming model for processing big data on a cluster

advantages:

- ▶ fault-tolerance
- ▶ elastic scaling
- ▶ integration with distributed file systems
- ▶ popular ecosystem — many good tools and services



Jeffrey Dean and Sanjay Ghemawat.

MapReduce: simplified data processing on large clusters. OSDI'04

Why translate automatically to MapReduce?

- ▶ although simple, MapReduce is not easy

Why translate automatically to MapReduce?

- ▶ although simple, MapReduce is not easy
- ▶ reduce cost of retargeting legacy imperative code

Why translate automatically to MapReduce?

- ▶ although simple, MapReduce is not easy
- ▶ reduce cost of retargeting legacy imperative code
- ▶ allow developers to concentrate on familiar imperative sequential code

WordCount

example: MapReduce WordCount


doc1

cat dog dog

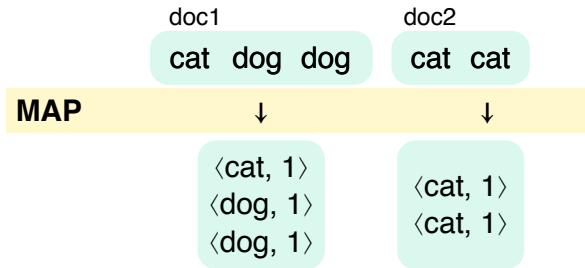
doc2

cat cat

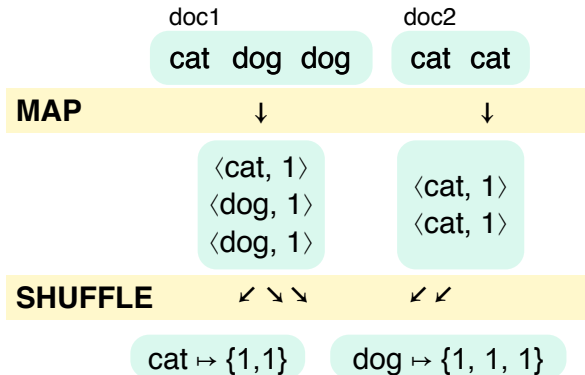
example: MapReduce WordCount



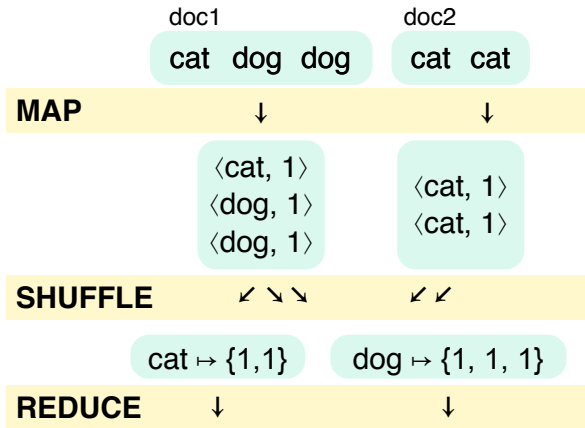
example: MapReduce WordCount



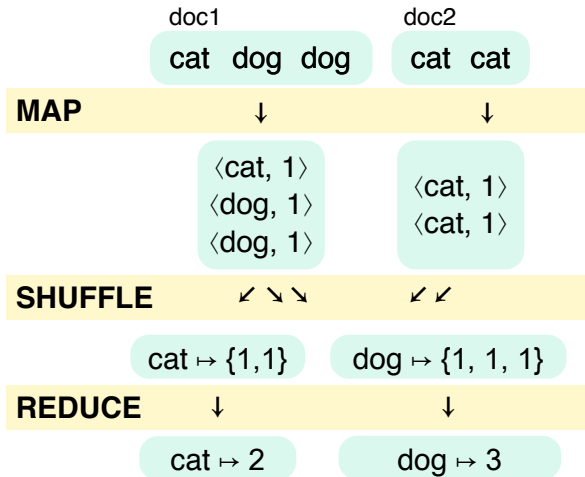
example: MapReduce WordCount



example: MapReduce WordCount



example: MapReduce WordCount



example

```
Map<String,Integer> count = new HashMap<>();

for (int i = 0; i < docs.size(); i++) {
    String[] words = tokenize(docs.get(i));
    for (int j = 0; j < words.length; j++) {
        String word = words[j];
        Integer prev = count.get(word);
        if (prev == null) prev = 0;
        count.put(word, prev + 1);
    }
}
```


example: imperative WordCount

```
Map<String,Integer> count = new HashMap<>();
```

```
for (int i = 0; i < docs.size(); i++) {  
    String[] words = tokenize(docs.get(i));  
    for (int j = 0; j < words.length; j++) {  
        String word = words[j];  
        Integer prev = count.get(word);  
        if (prev == null) prev = 0;  
        count.put(word, prev + 1);  
    }  
}
```

example: imperative WordCount

```
Map<String,Integer> count = new HashMap<>();
```

```
for (int i = 0; i < docs.size(); i++) {  
    String[] words = tokenize(docs.get(i));  
    for (int j = 0; j < words.length; j++) {  
        String word = words[j];  
        Integer prev = count.get(word);  
        if (prev == null) prev = 0;  
        count.put(word, prev + 1);  
    }  
}
```

example: imperative WordCount

```
Map<String,Integer> count = new HashMap<>();

for (int i = 0; i < docs.size(); i++) {
    String[] words = tokenize(docs.get(i));
    for (int j = 0; j < words.length; j++) {
        String word = words[j];
        Integer prev = count.get(word);
        if (prev == null) prev = 0;
        count.put(word, prev + 1);
    }
}
```

Example - imperative WordCount

```
Map<String,Integer> count = new HashMap<>();

for (int i = 0; i < docs.size(); i++) {
    String[] words = tokenize(docs.get(i));
    for (int j = 0; j < words.length; j++) {
        String word = words[j];
        Integer prev = count.get(word);
        if (prev == null) prev = 0;
        count.put(word, prev + 1);
    }
}
```

example: imperative WordCount

```
Map<String,Integer> count = new HashMap<>();

for (int i = 0; i < docs.size(); i++) {
    String[] words = tokenize(docs.get(i));
    for (int j = 0; j < words.length; j++) {
        String word = words[j];
        Integer prev = count.get(word);
        if (prev == null) prev = 0;
        count.put(word, prev + 1);
    }
}
```

example: imperative WordCount \Rightarrow MapReduce

```
Map<String,Integer> count = new HashMap<>();

for (int i = 0; i < docs.size(); i++) {
    String[] words = tokenize(docs.get(i));
    for (int j = 0; j < words.length; j++) {
        String word = words[j];
        Integer prev = count.get(word);
        if (prev == null) prev = 0;
        count.put(word, prev + 1);
    }
}
```



docs

```
.flatMap({ case (i, doc) => tokenize(doc) })
.map({ case (j, word) => (word, 1) })
.reduceByKey({ case (c1, c2) => c1 + c2 })
```

the MAP

```
Map<String,Integer> count = new HashMap<>();
```

```
for (int i = 0; i < docs.size(); i++) {  
    String[] words = tokenize(docs.get(i));  
    for (int j = 0; j < words.length; j++) {  
        String word = words[j];  
        Integer prev = count.get(word);  
        if (prev == null) prev = 0;  
        count.put(word, prev + 1);  
    }  
}
```



docs

```
.flatMap({ case (i, doc) => tokenize(doc) })  
.map({ case (j, word) => (word, 1) })  
.reduceByKey({ case (c1, c2) => c1 + c2 })
```

the SHUFFLE

```
Map<String,Integer> count = new HashMap<>();

for (int i = 0; i < docs.size(); i++) {
    String[] words = tokenize(docs.get(i));
    for (int j = 0; j < words.length; j++) {
        String word = words[j];
        Integer prev = count.get(word);
        if (prev == null) prev = 0;
        count.put(word, prev + 1);
    }
}
```



```
docs
  .flatMap({ case (i, doc) => tokenize(doc) })
  .map({ case (j, word) => (word, 1) })
  .reduceByKey({ case (c1, c2) => c1 + c2 })
```


the REDUCE

```
Map<String,Integer> count = new HashMap<>();
```

```
for (int i = 0; i < docs.size(); i++) {  
    String[] words = tokenize(docs.get(i));  
    for (int j = 0; j < words.length; j++) {  
        String word = words[j];  
        Integer prev = count.get(word);  
        if (prev == null) prev = 0;  
        count.put(word, prev + 1);  
    }  
}
```



docs

```
.flatMap({ case (i, doc) => tokenize(doc) })  
.map({ case (j, word) => (word, 1) })  
.reduceByKey({ case (c1, c2) => c1 + c2 })
```

MapReduce generated by MOLD

```
Map<String,Integer> count = new HashMap<>();
```

```
for (int i = 0; i < docs.size(); i++) {  
    String[] words = tokenize(docs.get(i));  
    for (int j = 0; j < words.length; j++) {  
        String word = words[j];  
        Integer prev = count.get(word);  
        if (prev == null) prev = 0;  
        count.put(word, prev + 1);  
    }  
}
```



docs

```
.flatMap({ case (i, doc) => tokenize(doc) })  
.map({ case (j, word) => (word, 1) })  
.reduceByKey({ case (c1, c2) => c1 + c2 })
```

MapReduce generated by MOLD

```
Map<String,Integer> count = new HashMap<>();


for (int i = 0; i < docs.size(); i++) {
    String[] words = tokenize(docs.get(i));
    for (int j = 0; j < words.length; j++) {
        String word = words[j];
        Integer prev = count.get(word);
        if (prev == null) prev = 0;
        count.put(word, prev + 1);
    }
}
```

⇓ **MOLD**


docs

```
.flatMap({ case (i, doc) => tokenize(doc) })
.map({ case (j, word) => (word, 1) })
.reduceByKey({ case (c1, c2) => c1 + c2 })
```

high-level “How?”

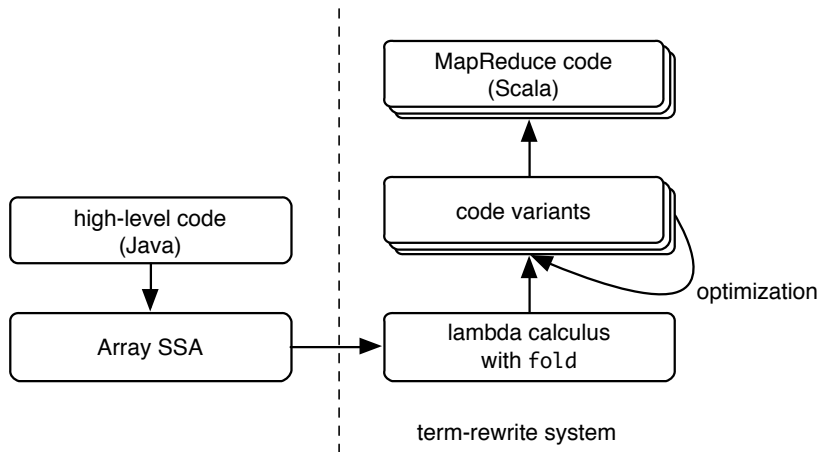


high-level “How?”



Kathleen Knobe and Vivek Sarkar. Array SSA form and its use in parallelization. POPL'98

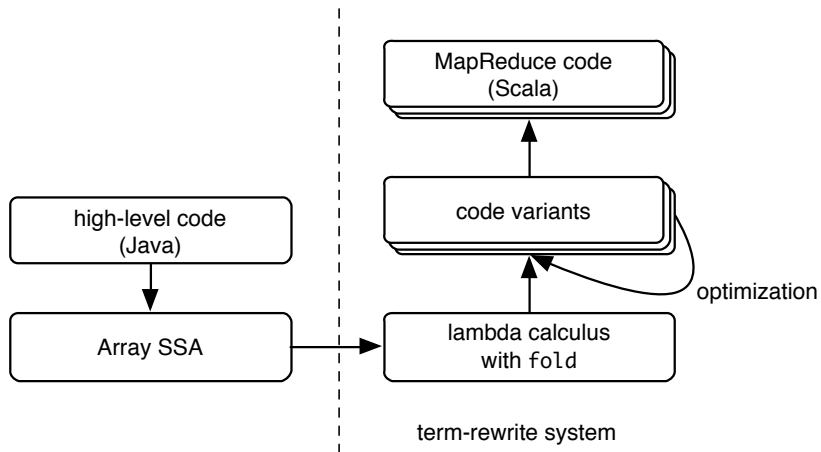
high-level “How?”



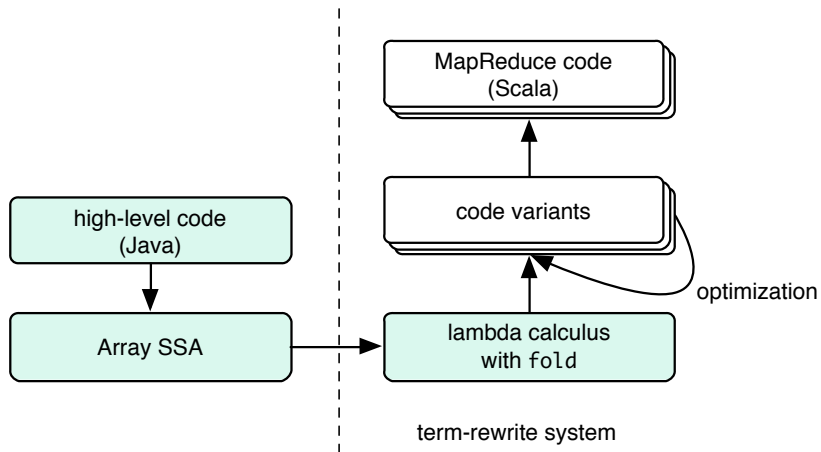
Andrew W. Appel. SSA is Functional Programming. '98

Richard Kelsey. A Correspondence between Continuation Passing Style and Static Single Assignment Form. '95

high-level “How?”



high-level “How?”



imperative \Rightarrow unoptimized functional

```
for (int i = 0; i < docs.size(); i++) {  
    String[] words = tokenize(docs.get(i));  
    for (int j = 0; j < words.length; j++) {  
        String word = words[j];  
        Integer prev = count.get(word);  
        if (prev == null) prev = 0;  
        count.put(word, prev + 1);  
    }  
}
```

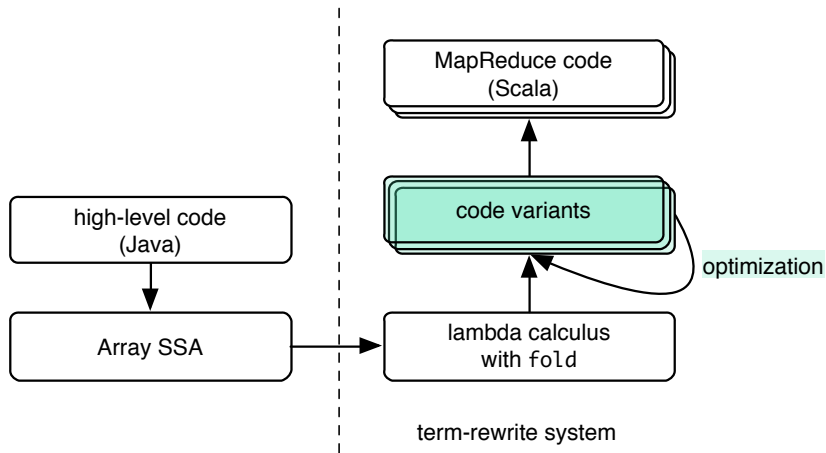


```
words.fold(count){ (count, word) =>  
    count.update(word, count(word) + 1)  
}
```

how to parallelize this?

```
words.fold(count){ (count, word) =>
    count.update(word, count(word) + 1)
}
```

high-level “How?”



trying a fold \Rightarrow map transformation ...

```
words.fold(count){ (count, word) =>
  count.update(word, count(word) + 1)
}
```

trying a fold \Rightarrow map transformation ...

```
words.fold(count){ (count, word) =>
  count.update(word, count(word) + 1)
}
```

...does not work:

distinct fold iterations can write to the same key in count

trying a fold \Rightarrow groupBy transformation ...

```
words.fold(count){ (count, word) =>
  count.update(word, count(word) + 1)
}
```

fold \Rightarrow groupBy

```
words.fold(count){ (count, word) =>  
  count.update(word, count(word) + 1)  
}
```



```
words.groupBy(word => word).map { (word, list) =>  
  list.fold(count(word)) { (sum, elem) => sum + 1 }  
}
```

fold \Rightarrow groupBy

```
words.fold(count){ (count, word) =>  
  count.update(word, count(word) + 1)  
}
```



```
words.groupBy(word => word).map { (word, list) =>  
  list.fold(count(word)) { (sum, elem) => sum + 1 }  
}
```

groupBy \equiv SHUFFLE

generic fold \Rightarrow groupBy

```
words.fold(count) { (count, word) =>
  count.update(word, count(word) + 1)
}
```

\Downarrow

```
words.groupBy(word => word).map { (word, list) =>
  list.fold(count(word)) { (sum, elem) => sum + 1 }
}
```

fold \Rightarrow groupBy

```
D.fold(A) { (a, k) => a.update(I, E) }
```

\Downarrow

```
D.groupBy( k => I ).map { (i, l) =>
  l.fold(A(i)) { (r, k) => E [ r / a(I) ] }}
```

$I \not\subseteq a$

$E \not\subseteq A(\neq I)$

transformation rules

fold \Rightarrow groupBy

$D.fold(A) \{ (a, k) \Rightarrow a.update(I, E) \}$

\Downarrow

$D.groupBy(k \Rightarrow I).map \{ (i, l) \Rightarrow$
 $l.fold(A(i)) \{ (r, k) \Rightarrow E [r / a(I)] \} \}$

$I \not\in a$

$E \not\in A(\neq I)$

transformation rules

fold \Rightarrow groupBy

$D.fold(A) \{ (a, k) \Rightarrow a.update(I, E) \}$
 \Downarrow
 $D.groupBy(k \Rightarrow I).map \{ (i, l) \Rightarrow l.fold(A(i)) \{ (r, k) \Rightarrow E [r / a(I)] \} \}$

$I \not\in a$
 $E \not\in A(\neq I)$

fold \Rightarrow map

$D.fold(A) \{ (a, k) \Rightarrow a.update(I, E) \}$
 \Downarrow
 $A.map \{ (k, v) \Rightarrow E [v / a(k)] \}$

$D = A.keys$
 $I = k$
 $E \not\in A(\neq I)$

transformation rules

fold \Rightarrow groupBy

$D.fold(A) \{ (a, k) \Rightarrow a.update(I, E) \}$
 \Downarrow
 $D.groupBy(k \Rightarrow I).map \{ (i, l) \Rightarrow$
 $l.fold(A(i)) \{ (r, k) \Rightarrow E [r / a(I)] \} \}$

$I \not\in a$
 $E \not\in A(\neq I)$

fold \Rightarrow map


$D.fold(A) \{ (a, k) \Rightarrow a.update(I, E) \}$
 \Downarrow
 $A.map \{ (k, v) \Rightarrow E [v / a(k)] \}$

$D = A.keys$
 $I = k$
 $E \not\in A(\neq I)$

... 16 more ...


Program variant exploration

at each step, MOLD can apply any of several transformation rules




Program variant exploration

the system is not confluent, nor terminating



Program variant exploration


MOLD attaches a cost to each program variant



$$\mathcal{C}(\text{map } Func) = \mathcal{C}_{init}^{\text{map}} + \mathcal{C}_{op}^{\text{map}} * \mathcal{C}(Func)$$


Program variant exploration

searches based on the cost (gradient descent)




Program variant exploration

another platform? different cost function



Program variant exploration

different cost function \Rightarrow different resulting programs



- ▶ applied MOLD on 7 programs (Phoenix benchmark suite¹)
 - ▶ WordCount
 - ▶ Image Histogram
 - ▶ LinearRegression
 - ▶ StringMatch
 - ▶ MatrixProduct
 - ▶ Principal Component Analysis (PCA)
 - ▶ K-Means

¹C. Ranger, R. Raghuraman, A. Penmetsa, G. Bradski, and C. Kozyrakis. Evaluating MapReduce for multi-core and multiprocessor systems. HPCA '07

evaluation methodology

Can MOLD generate *effective* MapReduce code?

Can MOLD generate *effective* MapReduce code?

- ▶ check semantics preservation

Can MOLD generate *effective* MapReduce code?

- ▶ check semantics preservation
- ▶ manually inspect the generated code
 - ▶ no redundant computation
 - ▶ high level of parallelism
 - ▶ accesses to large data structures should be localized

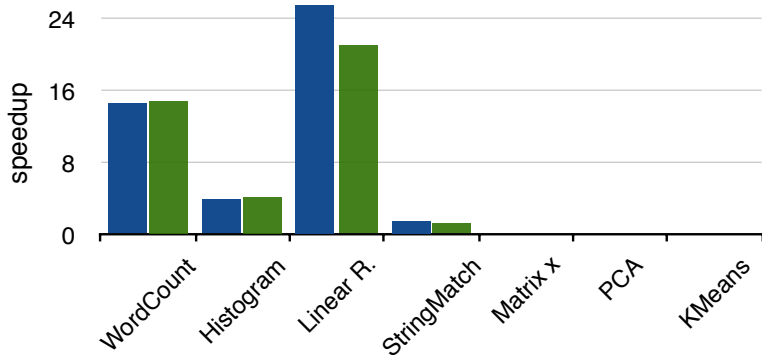
Can MOLD generate *effective* MapReduce code?

- ▶ no redundant computation — 5/7 programs
- ▶ parallelism — optimal for 4/7 programs
- ▶ memory accesses are localized — 5/7 programs

Can MOLD generate *effective* MapReduce code?

- ▶ check semantics preservation
- ▶ manually inspect the generated code
 - ▶ no redundant computation
 - ▶ high level of parallelism
 - ▶ accesses to large data structures should be localized
- ▶ execute with the three backends, compare with hand-written implementations:
 - ▶ Scala sequential collections
 - ▶ Scala parallel collections
 - ▶ Spark

comparison with hand-written implementations



- ▶ baseline: hand-written Scala using sequential collections
- ▶ blue is hand-written
- ▶ green is generated

Conclusions

- ▶ we propose an approach for transforming sequential imperative code to functional MapReduce
- ▶ sequential imperative \Rightarrow Array SSA \Rightarrow lambda with `fold`
- ▶ search through a space of possible optimizations
 - ▶ transformations are expressed as rewrite rules
 - ▶ generic handling of indirect updates
 - ▶ cost function can be platform-dependent
- ▶ good results on a small set of benchmarks

Appendix

Cost estimation

$$\mathcal{C}(F \circ G) = \mathcal{C}(F) + \mathcal{C}(G)$$

$$\mathcal{C}(F(G)) = \mathcal{C}(F) + \mathcal{C}(G)$$

$$\mathcal{C}(\langle F, G, \dots \rangle) = \mathcal{C}(F) + \mathcal{C}(G) + \dots$$

$$\mathcal{C}(A[I]) = C_{get}^{collection} + \mathcal{C}(A) + \mathcal{C}(I)$$

$$\mathcal{C}(A[K := V]) = C_{set}^{collection} + \mathcal{C}(A) + \mathcal{C}(K) + \mathcal{C}(V)$$

$$\mathcal{C}(\text{map } F) = C_{init}^{\text{map}} + C_{op}^{\text{map}} * \mathcal{C}(F)$$

$$\mathcal{C}(\text{fold } I \ F) = \mathcal{C}(I) + C_{init}^{\text{fold}} + C_{op}^{\text{fold}} * \mathcal{C}(F)$$

$$\mathcal{C}(\text{groupBy } F) = C_{init}^{\text{groupBy}} + C_{op}^{\text{groupBy}} * \mathcal{C}(F)$$

Transformation rules

(extract map from fold)

$$\frac{\text{fold}\langle r_0^0, \dots, r_n^0 \rangle \lambda \langle r_0, \dots, r_n \rangle K V . E}{(\text{fold}\langle r_0^0, \dots, r_n^0 \rangle \lambda \langle r_0, \dots, r_n \rangle K \langle v_0^f, \dots, v_m^f \rangle V_{\cap \text{free}(F)} . F) \circ (\text{map} \lambda K V . \langle G[r_-^0/r_-], V_{\cap \text{free}(F)} \rangle)}$$

(fold to group by)

$$\frac{\text{fold } r_0 \lambda r V . r[E := B]}{(\text{map} \lambda k l . (\text{fold } r_0[k] \lambda g V . C) l) \circ (\text{groupBy} \lambda V . E)}$$

$$E = (\lambda \langle v_0^f, \dots, v_m^f \rangle . F) \circ G$$

F is $\arg \max \mathcal{C}(G)$ with the condition:

$$\nexists i \in [0..n] . r_i \in G \wedge r_i \in E[r_-^0/r_-]$$

where

$$r_-^0/r_- = r_i^0[k]/r_i[k] \text{ applied for all } i \in [1..n] \ k \in K$$

$$C = B[g/r[E]]$$

$$r \notin C \wedge r \notin E \wedge \exists v \in V . v \in E$$

we cannot prove E is distinct across the folding

Is the proposed approach general?

algorithm	loops/ loop nests	translation time (s)	transformations
WordCount	2/1	11	15
Histogram	1/1	233	18
LinearRegression	1/1	28	2
StringMatch	1/1	68	2
Matrix \times	3/1	40	20
PCA	5/2	66	15
KMeans	6/2	340	10