Hasso Plattner Institut

IT Systems Engineering | Universität Potsdam

Link Analysis goes to the Stratosphere

Thomas Zimmermann
Philipp Berger

Recap

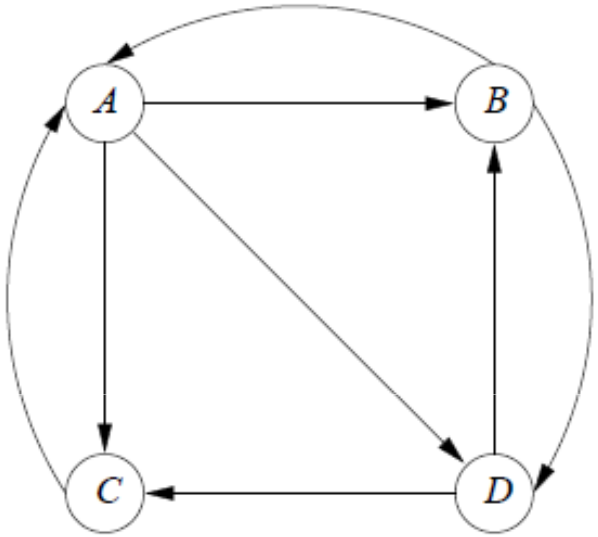

$$\begin{pmatrix} m_{1,1} & m_{1,2} & m_{1,3} & m_{1,4} \\ m_{2,1} & m_{2,2} & m_{2,3} & m_{2,4} \\ m_{3,1} & m_{3,2} & m_{3,3} & m_{3,4} \\ m_{4,1} & m_{4,2} & m_{4,3} & m_{4,4} \end{pmatrix} \begin{pmatrix} v_1 \\ v_2 \\ v_3 \\ v_4 \end{pmatrix}$$

Agenda

1. PageRank on Stratosphere
2. Hadoop vs. Stratosphere
3. Evaluation

Agenda

1. **PageRank on Stratosphere**
2. Hadoop vs. Stratosphere
3. Evaluation

PageRankPlan

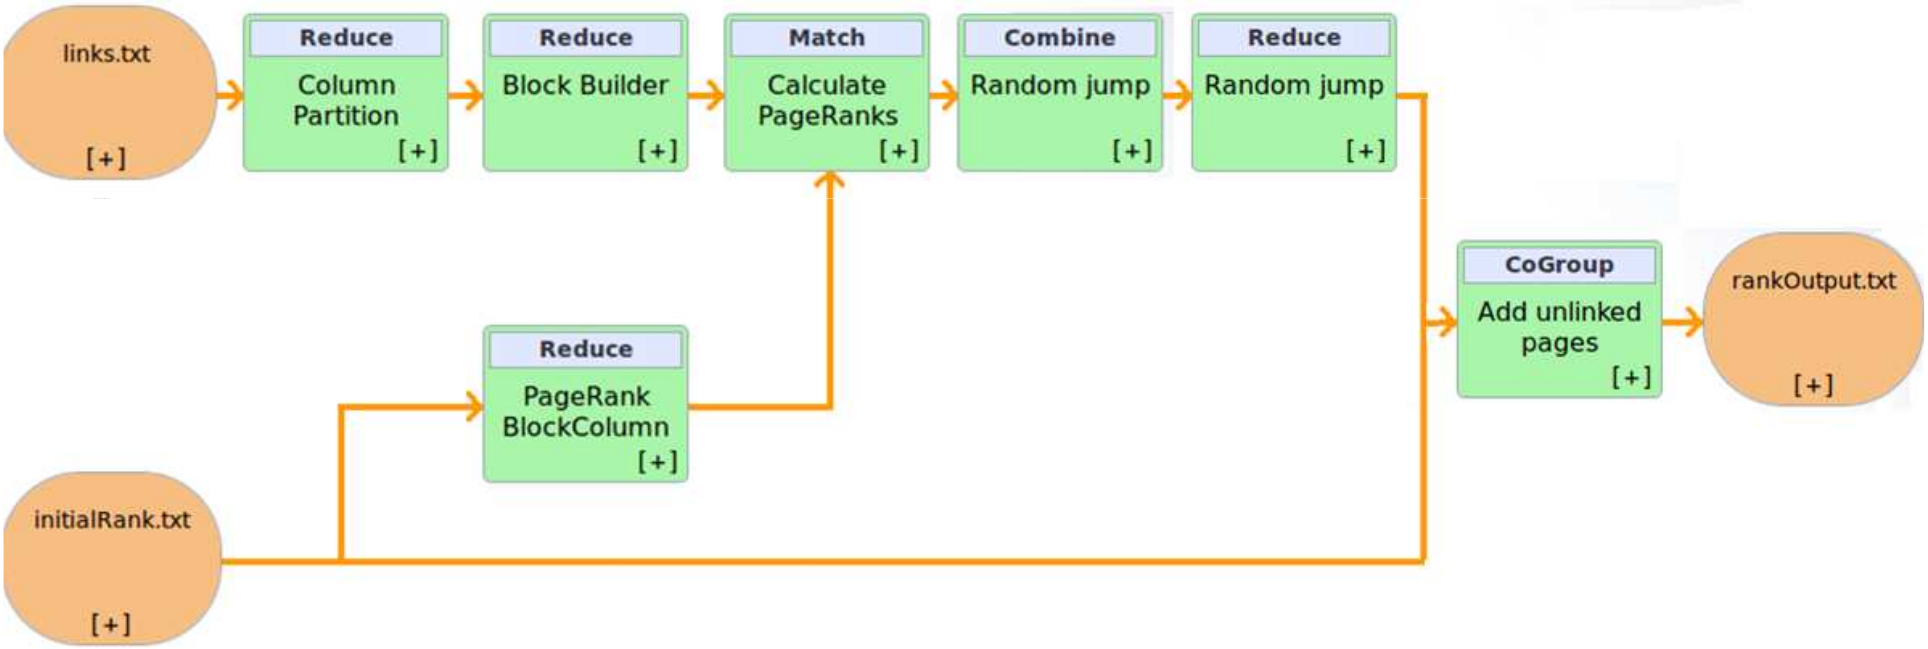

PageRankPlan

- Block Building

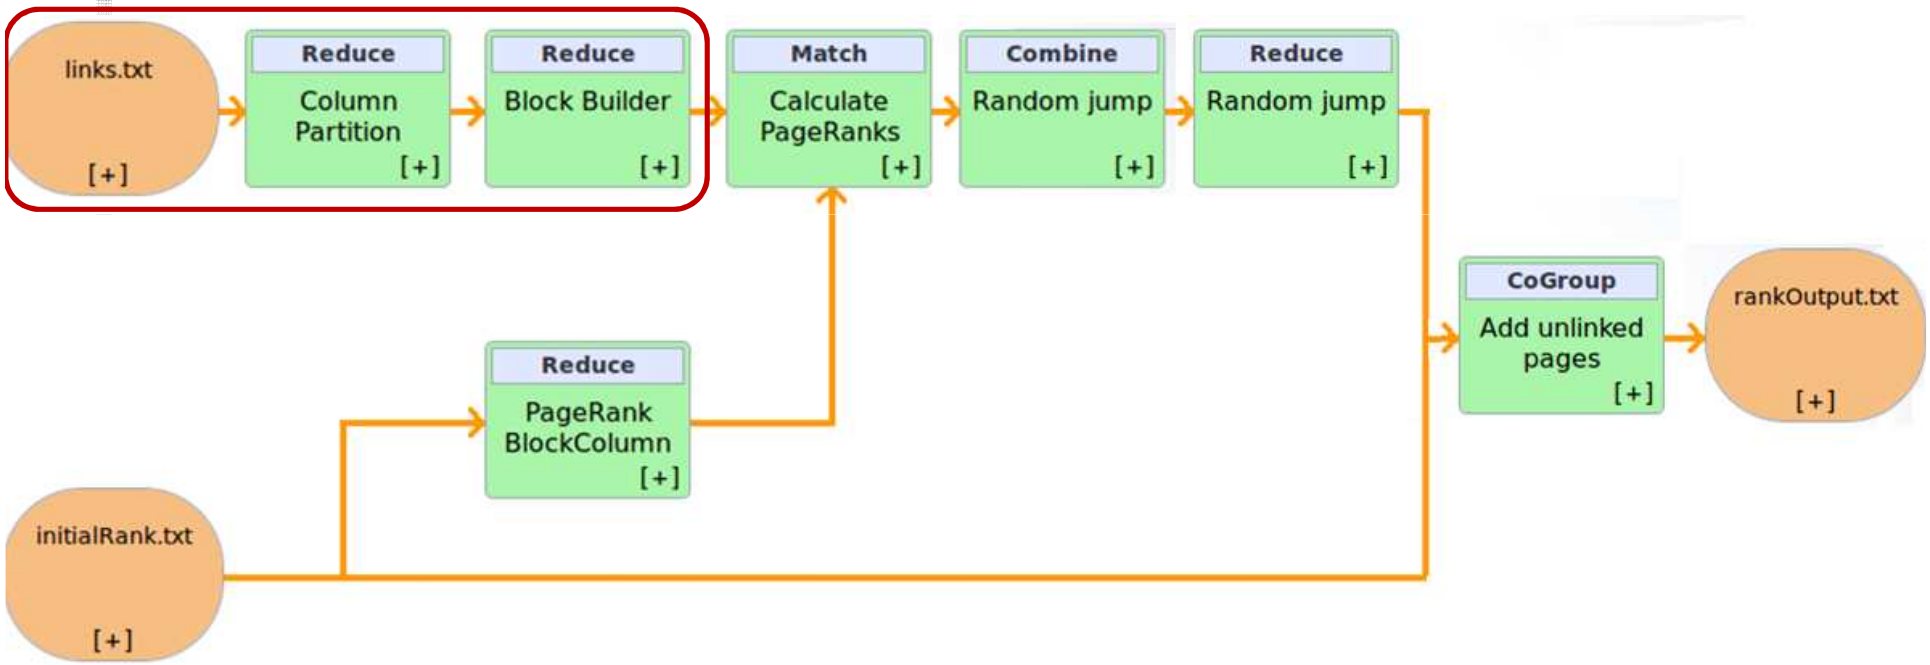

InputFormat Links

Innsbruck	
Country	Austria
State	Tyrol
Administrative region	Statutory city
Population	117,342 (2006)
Area	104.91 km²
Population density	1,119 /km²
Elevation	574 m
Coordinates	47°16' N 11°23' E
Postal code	6010-6080
Area code	0512
Licence plate code	I
Mayor	Hilde Zach
Website	www.innsbruck.at

Record
Source Page Id
Target Page Id

BlockBuilding

- Column Partitioner Reducer

	$32_{\text{Eiffel Tower}}$
122_{France}	X
123_{Japan}	
124_{Paris}	X

Record
Block Column Id
Block Row Id
BlockColumnValue

```
class BlockColumnValue implements Value


 long sourceSiteId;
 Collection<Long> targetSiteIds;
 long targetSetSize;
```

BlockBuilding

- Column Partitioner Reducer

	$32_{\text{Eiffel Tower}}$
122_{France}	X
123_{Japan}	
124_{Paris}	X

Record	
Block Column Id	16
Block Row Id	61
BlockColumnValue	(32, [122], 2)

```
class BlockColumnValue implements Value


 long sourceSiteId;
 Collection<Long> targetSiteIds;
 long targetSetSize;
```

BlockBuilding

- BlockPartConcatenate

	$32_{\text{Eiffel Tower}}$	33_{Sushi}
122_{France}	X	
123_{Japan}		X
124_{Paris}	X	

Record
Block Column Id
Block Row Id
BlockColumnValueList

```
class BlockColumnValueList extends PactList<BlockColumnValue>
```

BlockBuilding

- BlockPartConcatenate

	$32_{\text{Eiffel Tower}}$	33_{Sushi}
122_{France}	X	
123_{Japan}		X
124_{Paris}	X	

Record	
Block Column Id	16
Block Row Id	61
BlockColumnValueList	(32, [122], 2), (33, [123], 1)

```
class BlockColumnValueList extends PactList<BlockColumnValue>
```

PageRankPlan

- Read Input Ranks

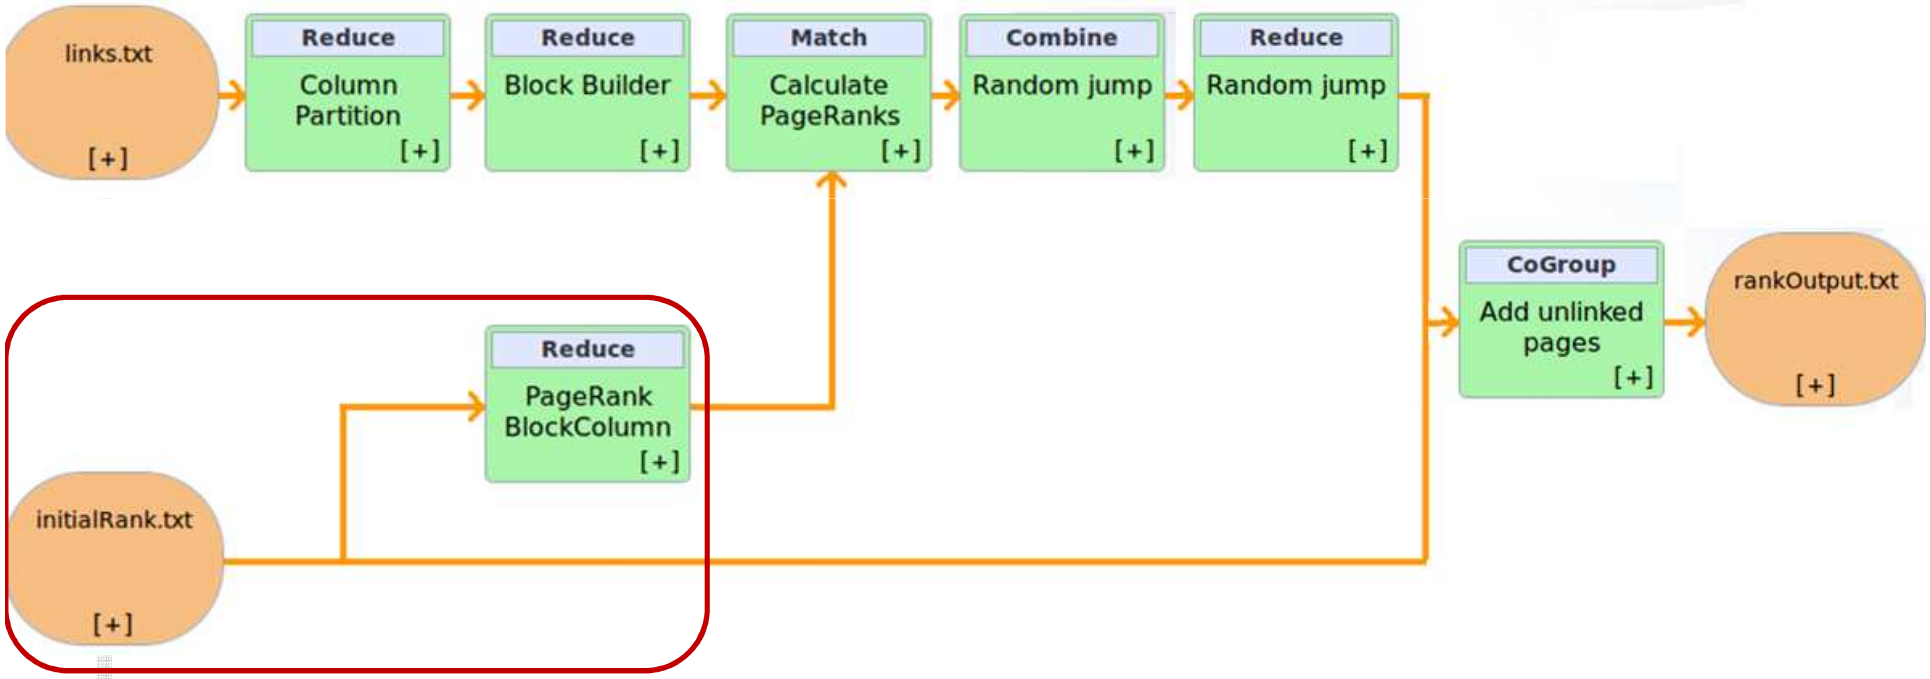

Inputformat Ranks

- Assign block column id to pages
  - To group page ranks for each block

Page	Rank
$32_{\text{Eiffel Tower}}$	0.125
33_{Sushi}	0.125
34_{Island}	0.125

Record
Block Column Id
Page Id
Rank

Inputformat Ranks

- Assign block column id to pages
  - To group page ranks for each block

Page	Rank
$32_{\text{Eiffel Tower}}$	0.125
33_{Sushi}	0.125
34_{Island}	0.125

Record	
Block Column Id	16
Page Id	33
Rank	0.125

RankColumn Reducer

- Creates Map for each Block

Record	
Block Column Id	16
Page Id	32
Rank	0.125

Record	
Block Column Id	16
Page Id	33
Rank	0.125

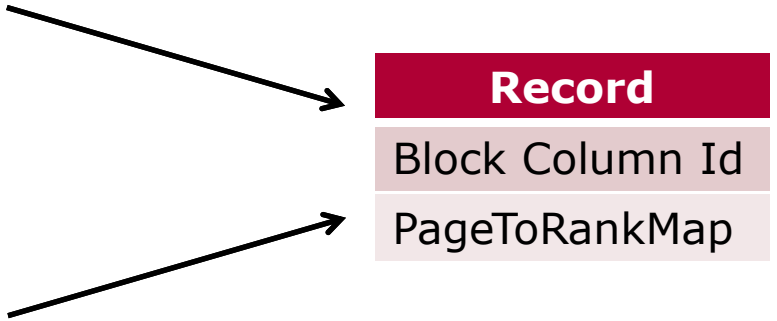

Record
Block Column Id
PageToRankMap

```
class PageToRankMap extends PactMap<PactLong, PactDouble>
```

RankColumn Reducer

- Creates Map for each Block

Record	
Block Column Id	16
Page Id	32
Rank	0.125

Record	
Block Column Id	16
Page Id	33
Rank	0.125

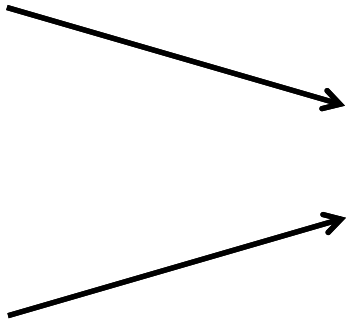

Record	
Block Column Id	16
PageToRankMap	32 → 0.125 33 → 0.125

```
class PageToRankMap extends PactMap<PactLong, PactDouble>
```

PageRankPlan

- Calculate PageRank

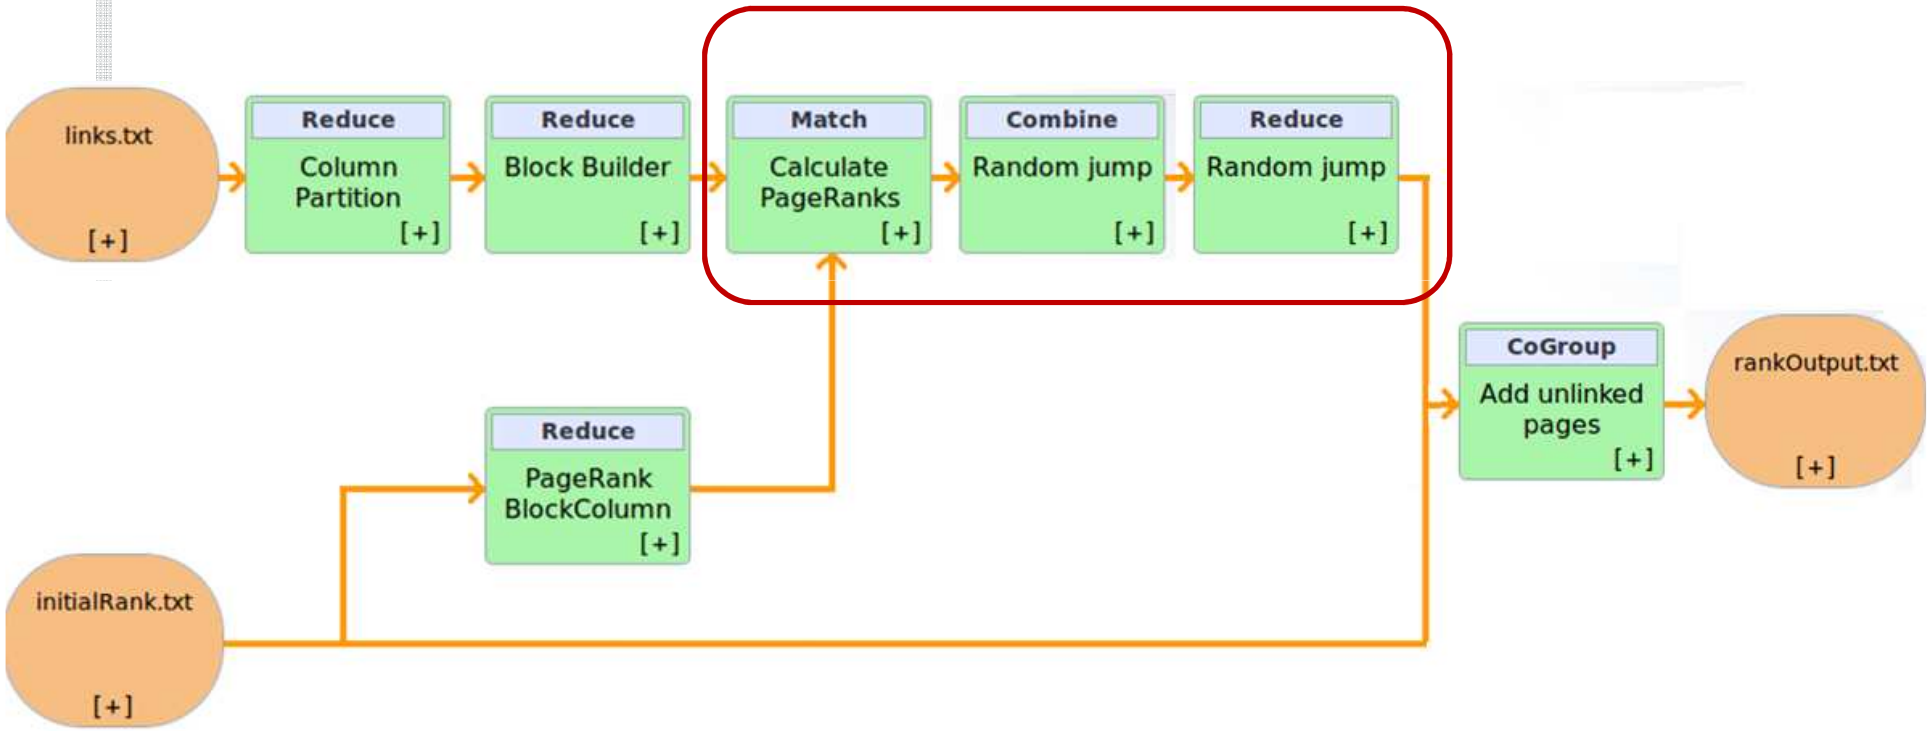

Stratosphere Match

- Like a database JOIN

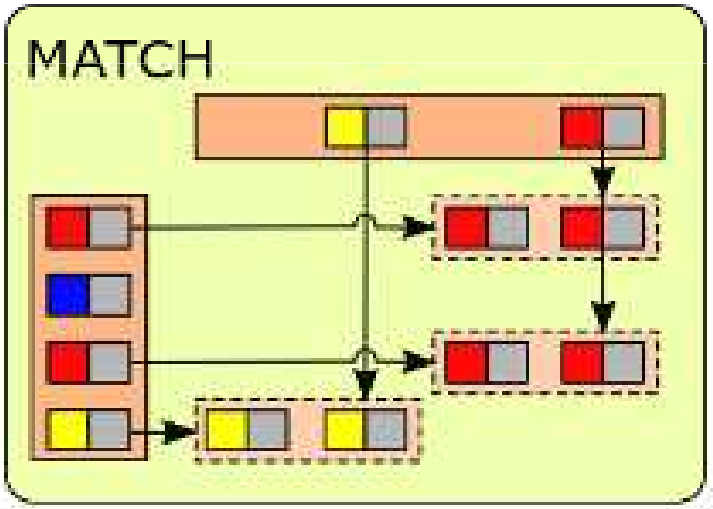

Match Rank and Block

PageToRankMap	
Block Column Id	16
PageToRankMap	32 → 0.125 33 → 0.125

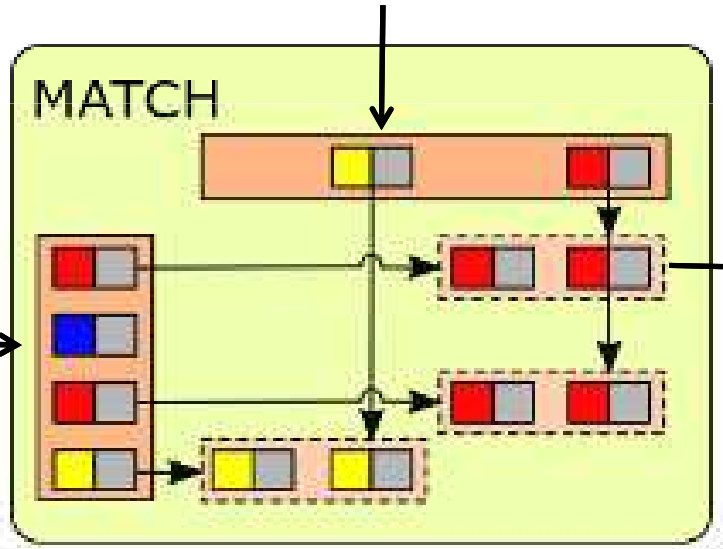

Record	
Block Column Id	16
Block Row Id	61
BlockColumn ValueList	(32,[122],2), (33,[123],1)

Record
Target Page Id
Partial Rank

Match Rank and Block

PageToRankMap	
Block Column Id	16
PageToRankMap	32 → 0.125 33 → 0.125

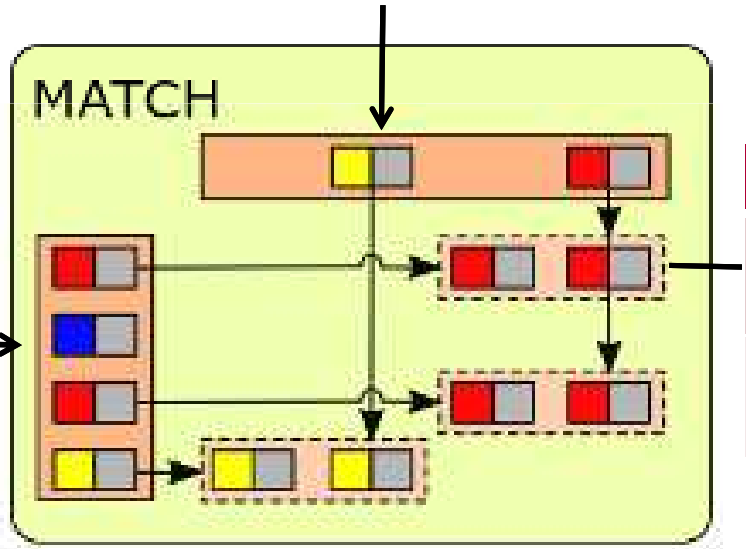

Record	
Block Column Id	16
Block Row Id	61
BlockColumn ValueList	(32,[122],2), (33,[123],1)

Record	
Target Page Id	122
Partial Rank	0.0625

Reduce to get final Rank

- RandomJump
  - Combine
    - Sums Up new ranks
  - Reduce
    - Calculates the final sum
    - Adds the RandomJump component

Record
Target Page Id
Calculated Rank

```
// Random Jump
double result = beta * pageRankSum + (1 - beta) * (1 / (double)pageCount);
```

PageRankPlan

- Calculate PageRank

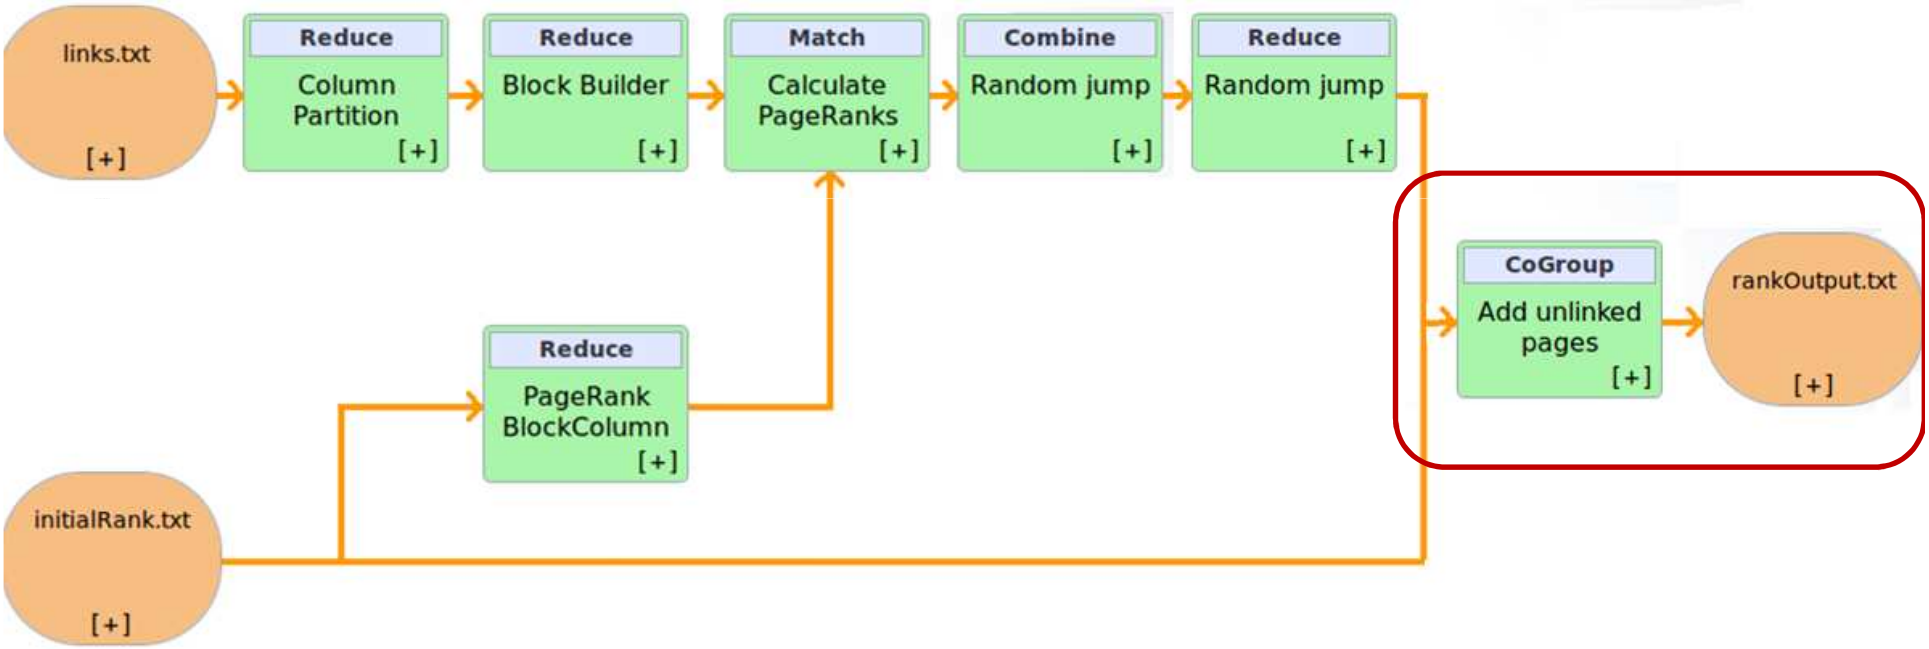

Some ranks are missing?

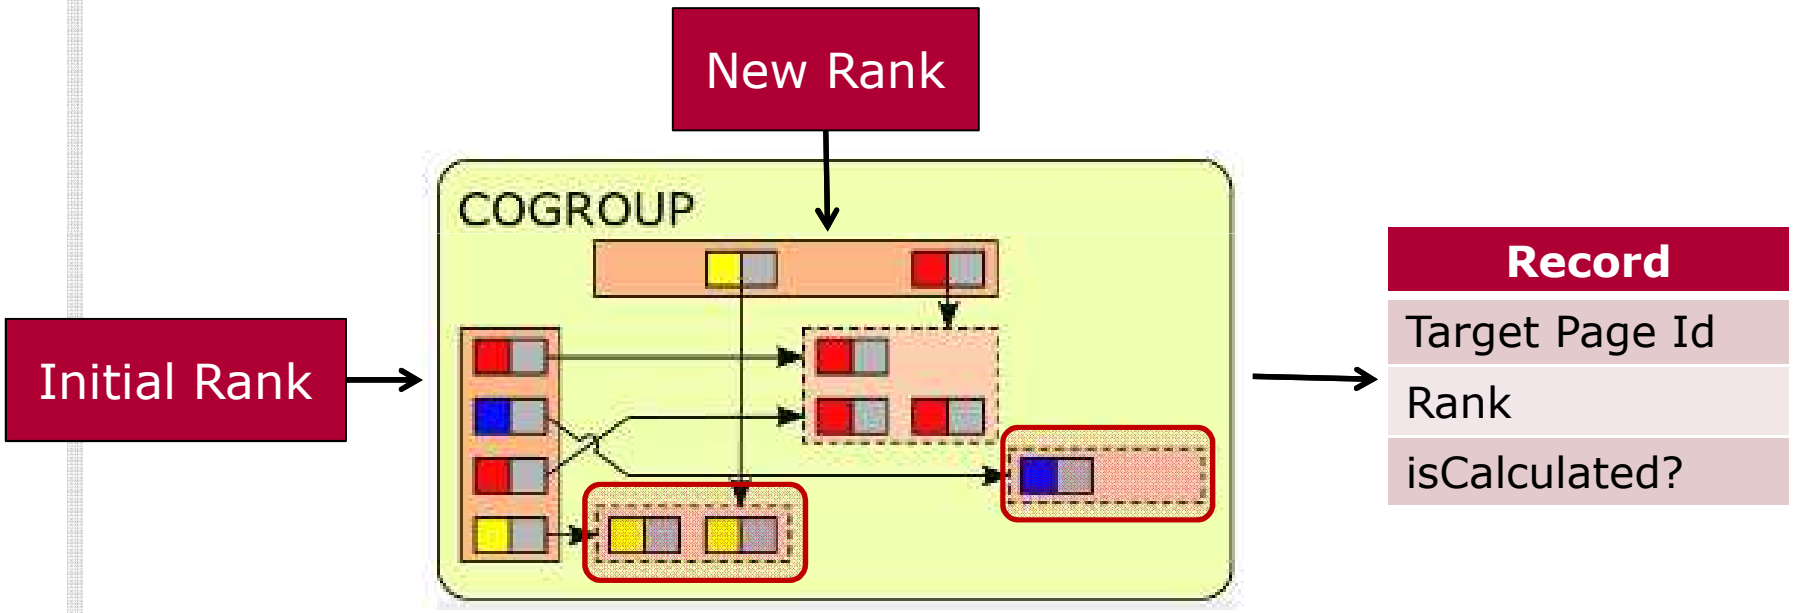

- $(InitialRanks \backslash NewRanks) \cup NewRanks$

PageRankPlan

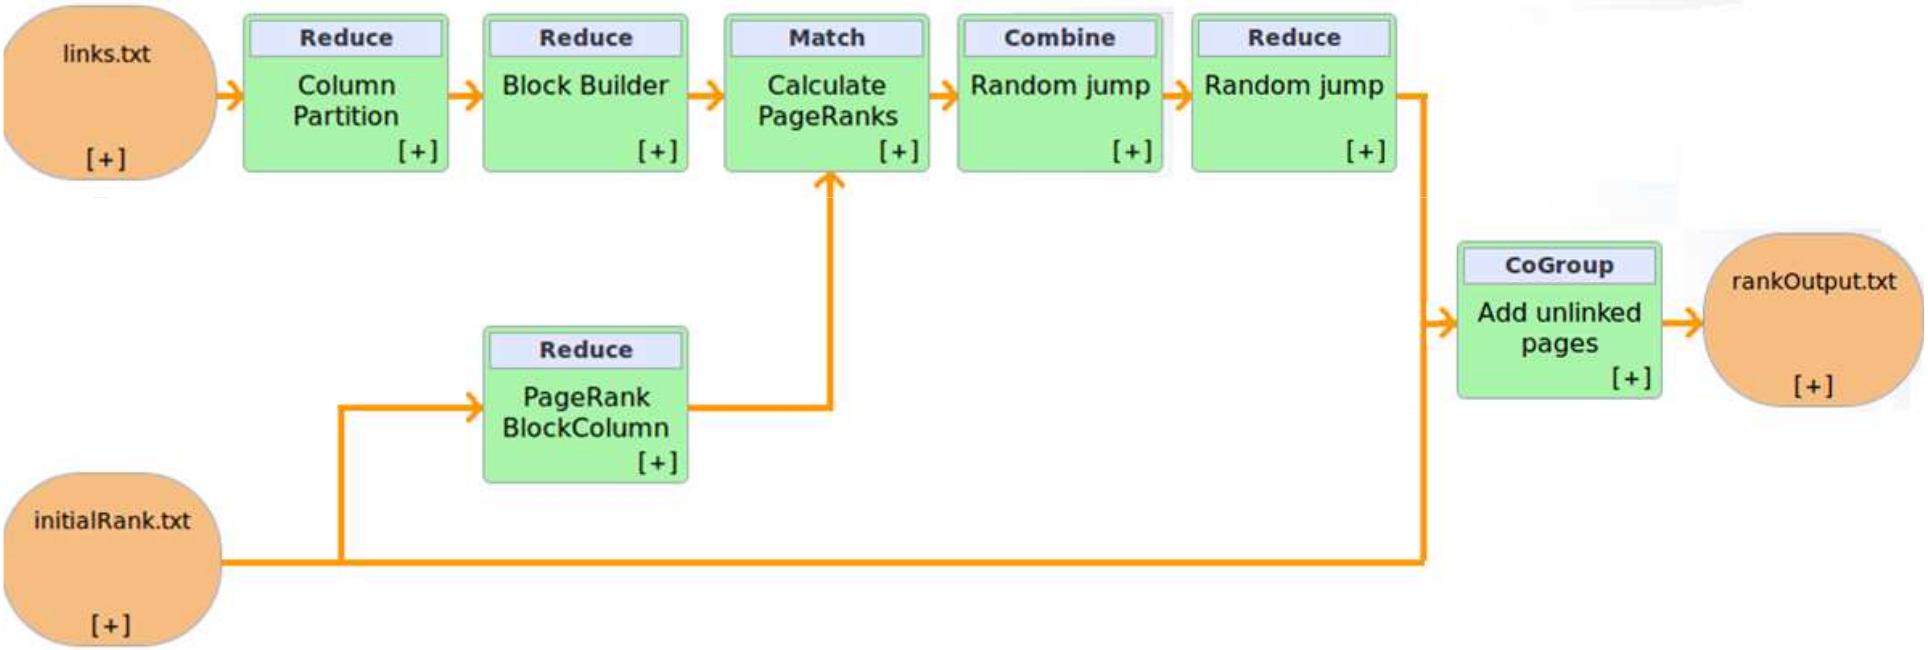

Does it converge?

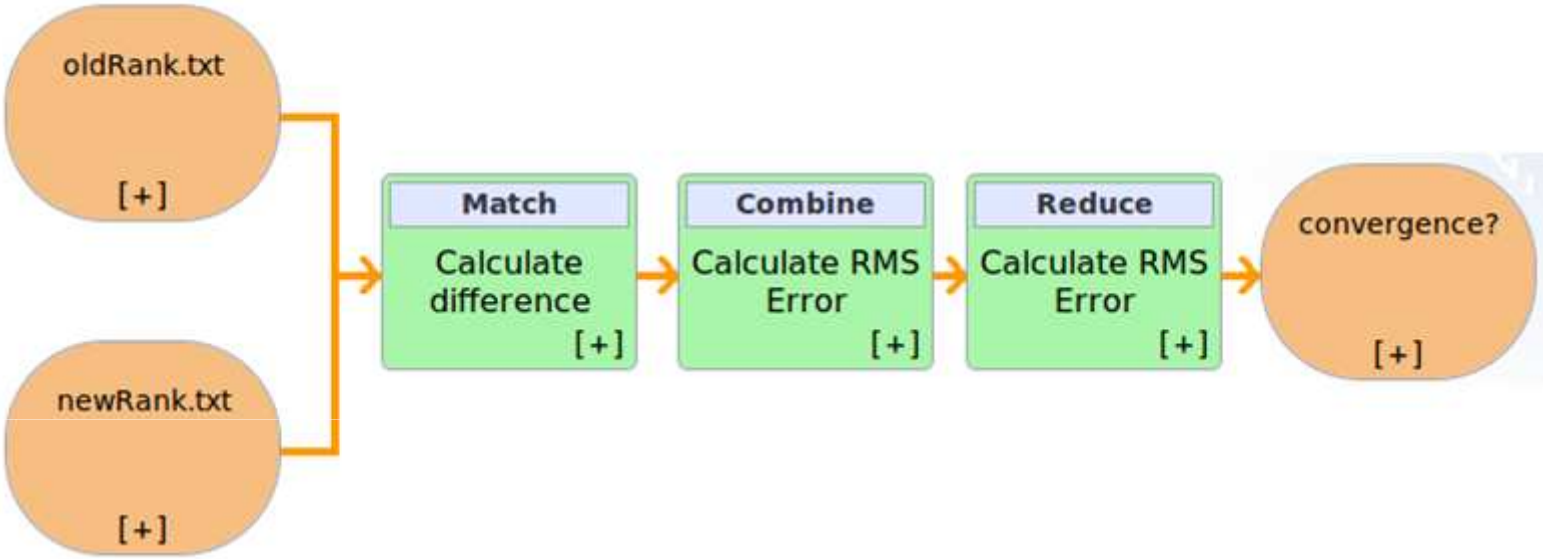

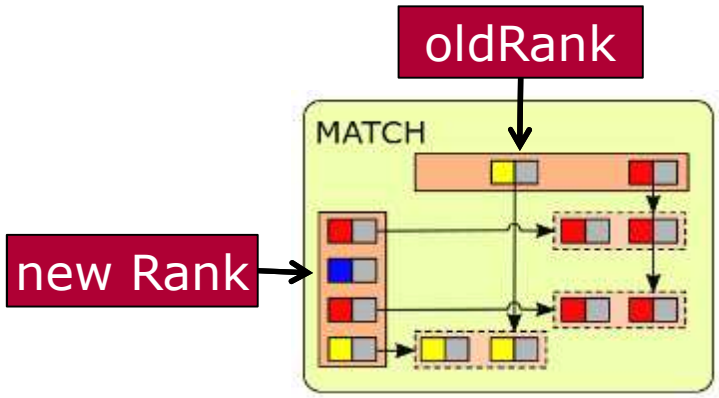

Result Ranking

Rank	PageRank	Site
1	0.0097279421588840 3	<http://www.opengis.net/gml/_**Feature**>
2	0.0046402352314071 56	<http://dbpedia.org/resource/**United_States**>
3	0.0033554866451856414	<http://dbpedia.org/resource/**Eukaryote**>
4	0.0019618452796001246	<http://dbpedia.org/resource/**Animal**>
5	0.0013024558125835565	<http://dbpedia.org/resource/**France**>
6	0.0011861730946363563	<http://dbpedia.org/resource/**United_Kingdom**>
7	0.001178379282062479	<http://dbpedia.org/resource/**English_language**>
8	0.0011006482651497 23	<http://dbpedia.org/resource/**Plant**>
9	0.001082222684959 4	<http://dbpedia.org/resource/**Poland**>
10	0.0008418068734001761	<http://dbpedia.org/resource/**Australia**>

Agenda

1. PageRank on Stratosphere
2. **Hadoop vs. Stratosphere**
3. Evaluation

Hadoop vs. Stratosphere

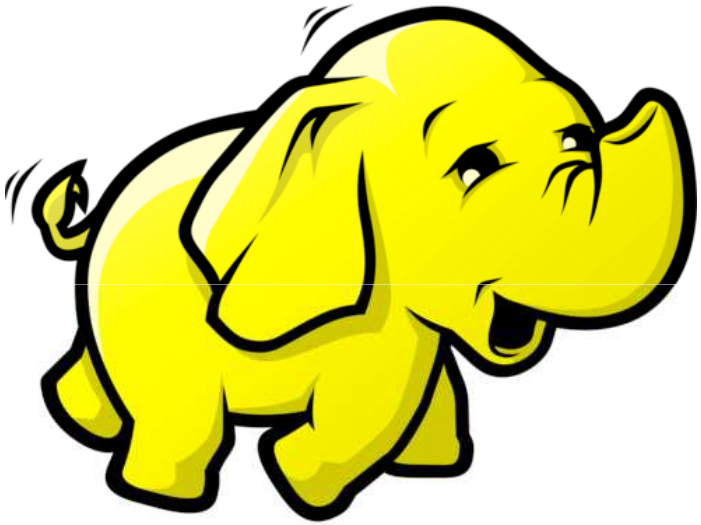

VS

- *from the coders point of view*

InputFormat and Mapper

Hadoop

- Need for an extra mapper for input file splitting
- Complex structure of InputFormat in combination with RecordReader and InputSplit

Stratosphere

- Easy adaptable InputFormat interface
- Definition of keys in the contract

Type Safety

- Hadoop
  - Types of key and value at compile time via Generics
  - TupleWritable
- Stratosphere
  - Types of record fields only known at runtime
  - Combined keys

Distributed Cache and Match

Hadoop

- Distributed Cache spreads source ranks to each mapper
- Each mapper reads from same Cache

Stratosphere

- Combine multiple inputs with Match
- Match gets the subset of ranks for one block

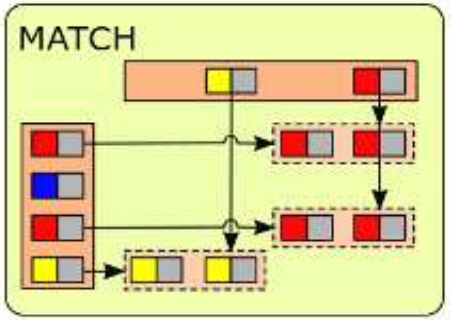

Need for Initial Rank

Stratosphere with Match

- explicit initial rank
- Co-Group for unlinked pages

Hadoop with Distributed Cache

- On-the-fly default page rank calculation

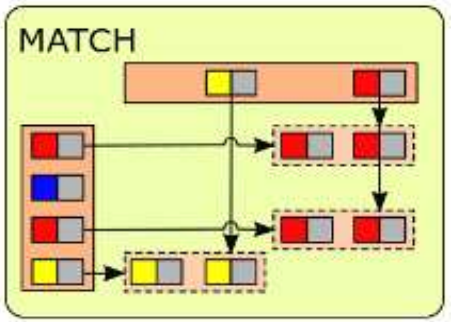

Stability

- Hadoop
  - Windows support
  - Jobtracker
- Stratosphere
  - Explicit degree of parallelism
  - Hangs sometimes without error

Agenda

1. PageRank on Stratosphere
2. Hadoop vs. Stratosphere
3. **Evaluation**

Convergence

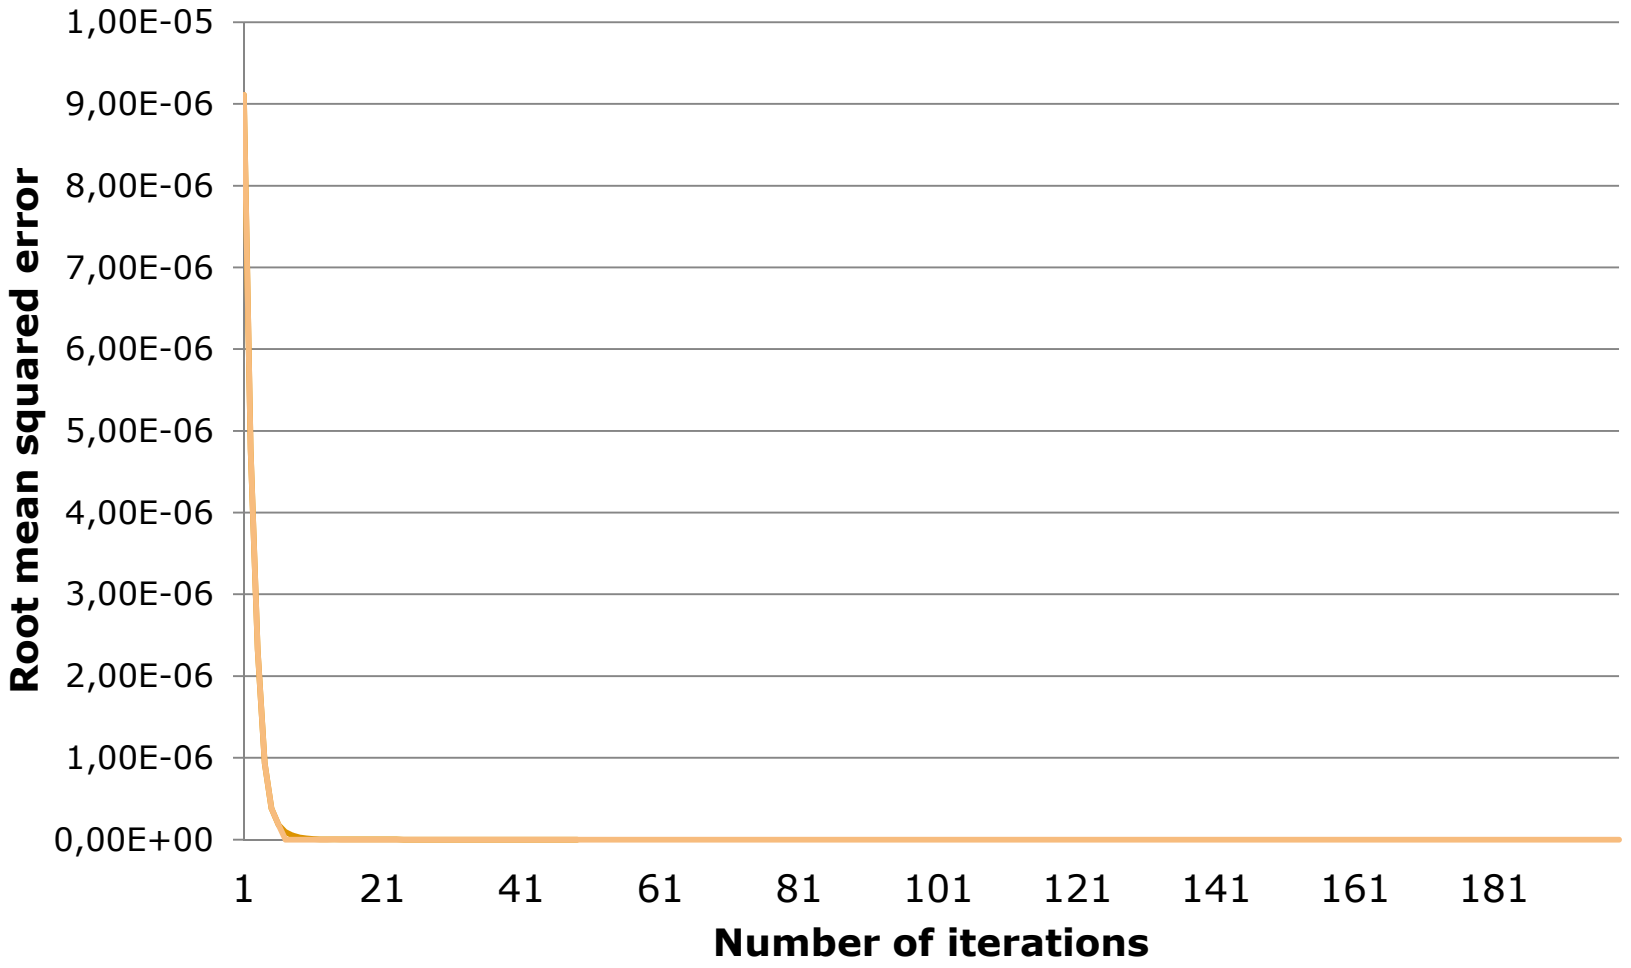

Convergence

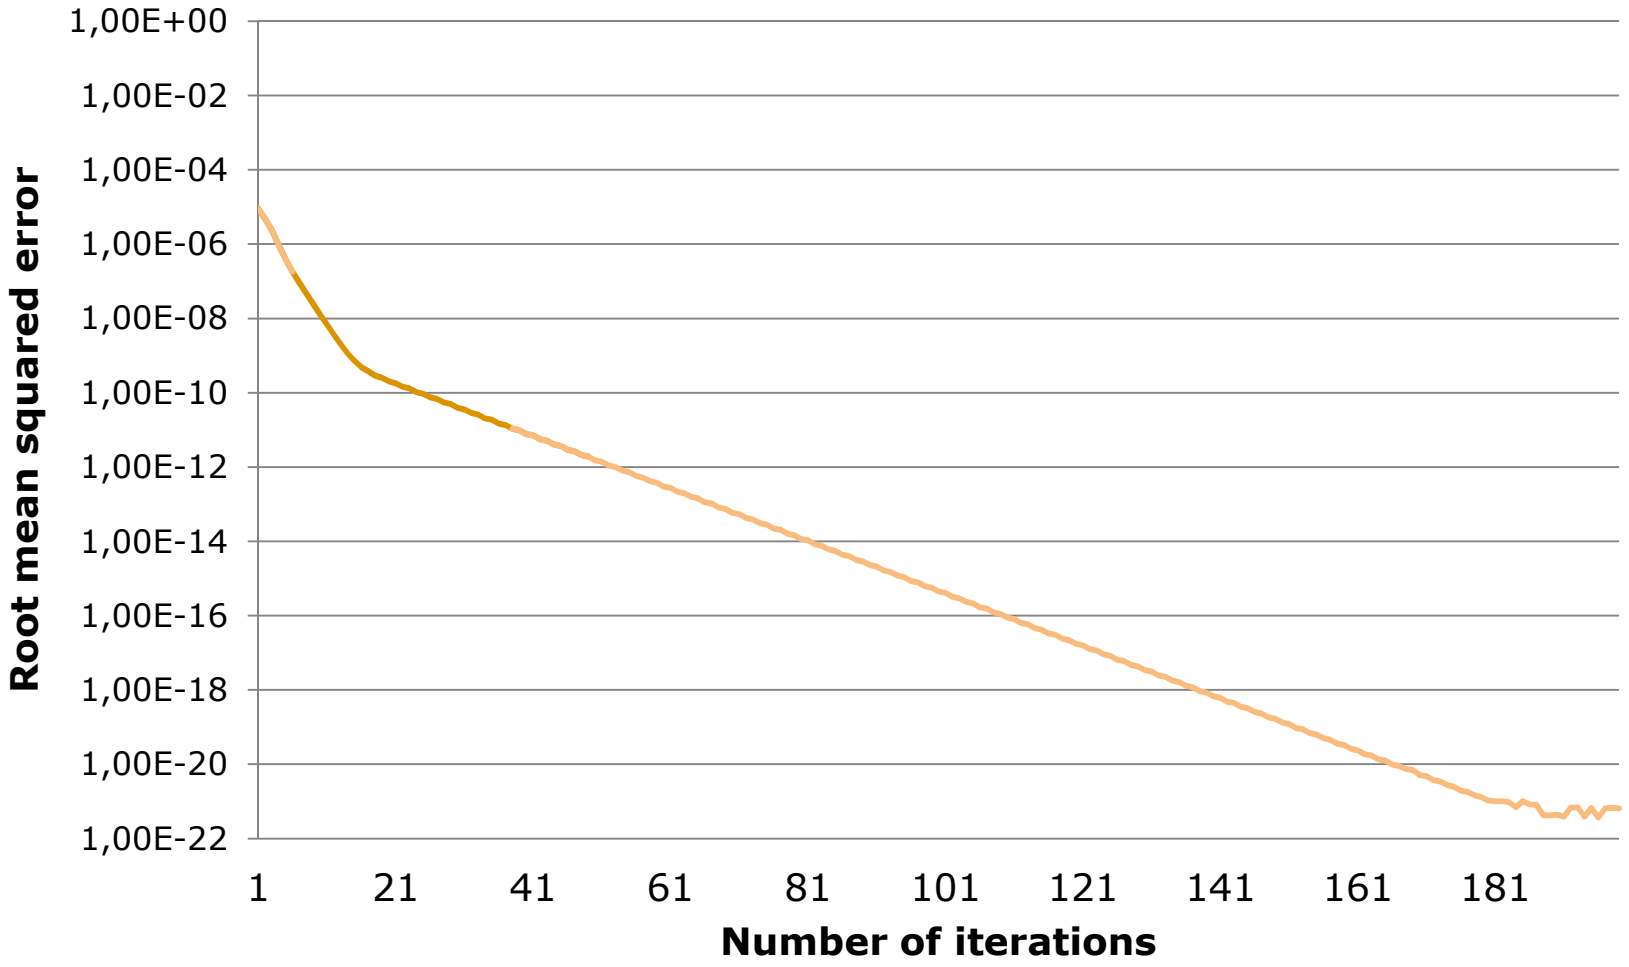

Runtime with different input sizes

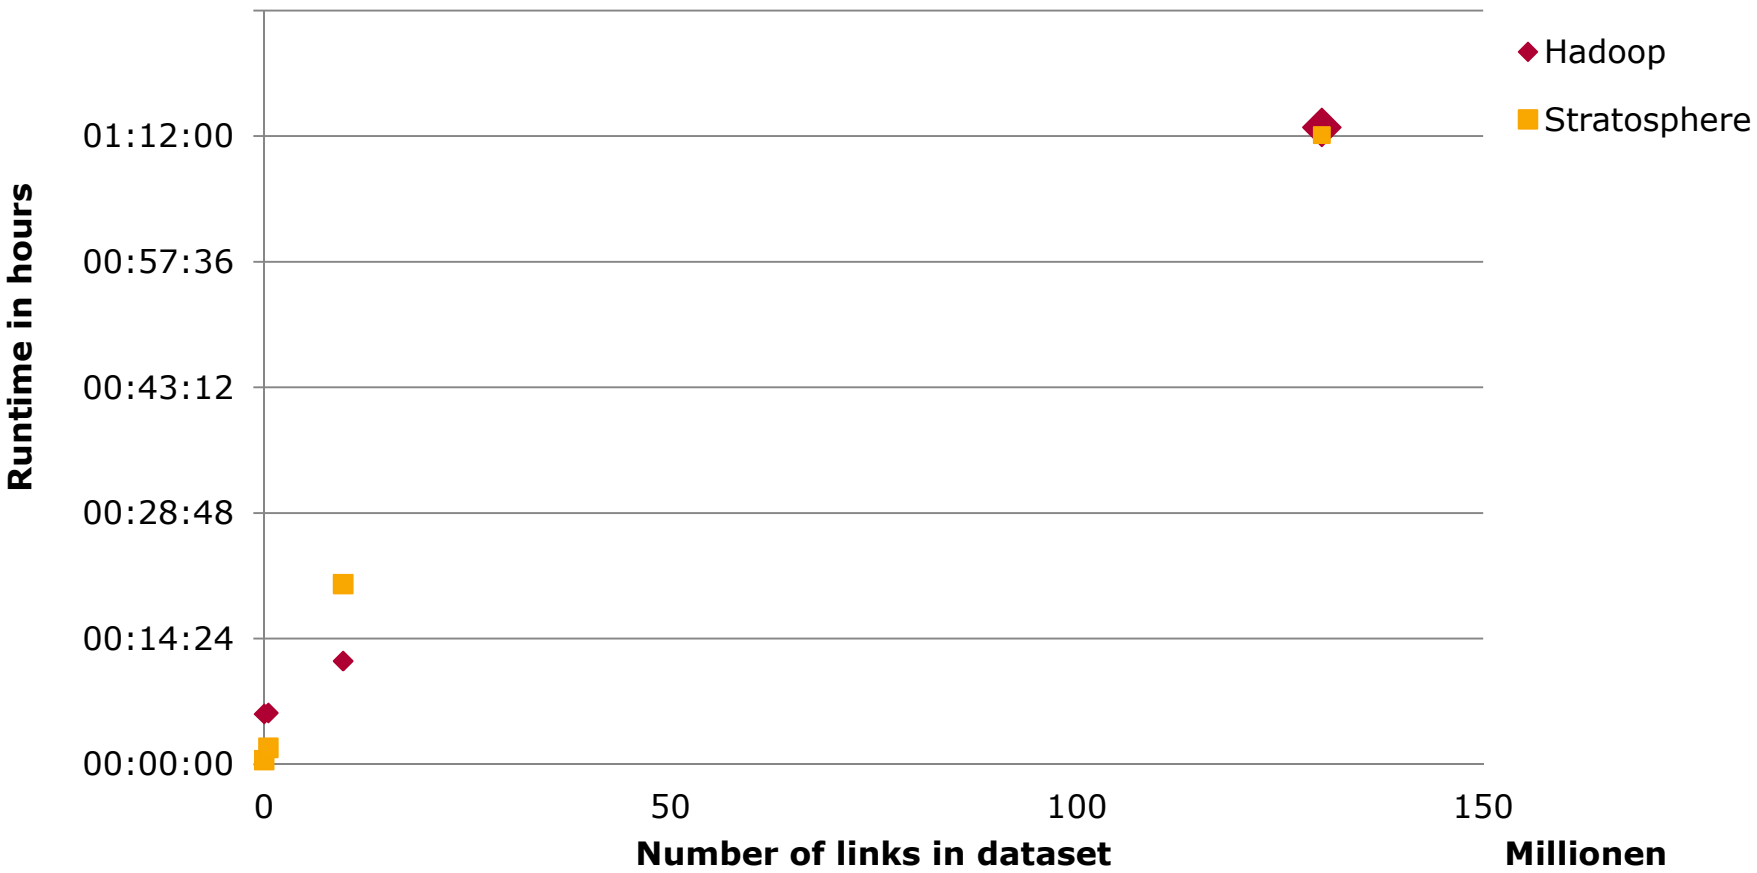

Runtime with different block sizes

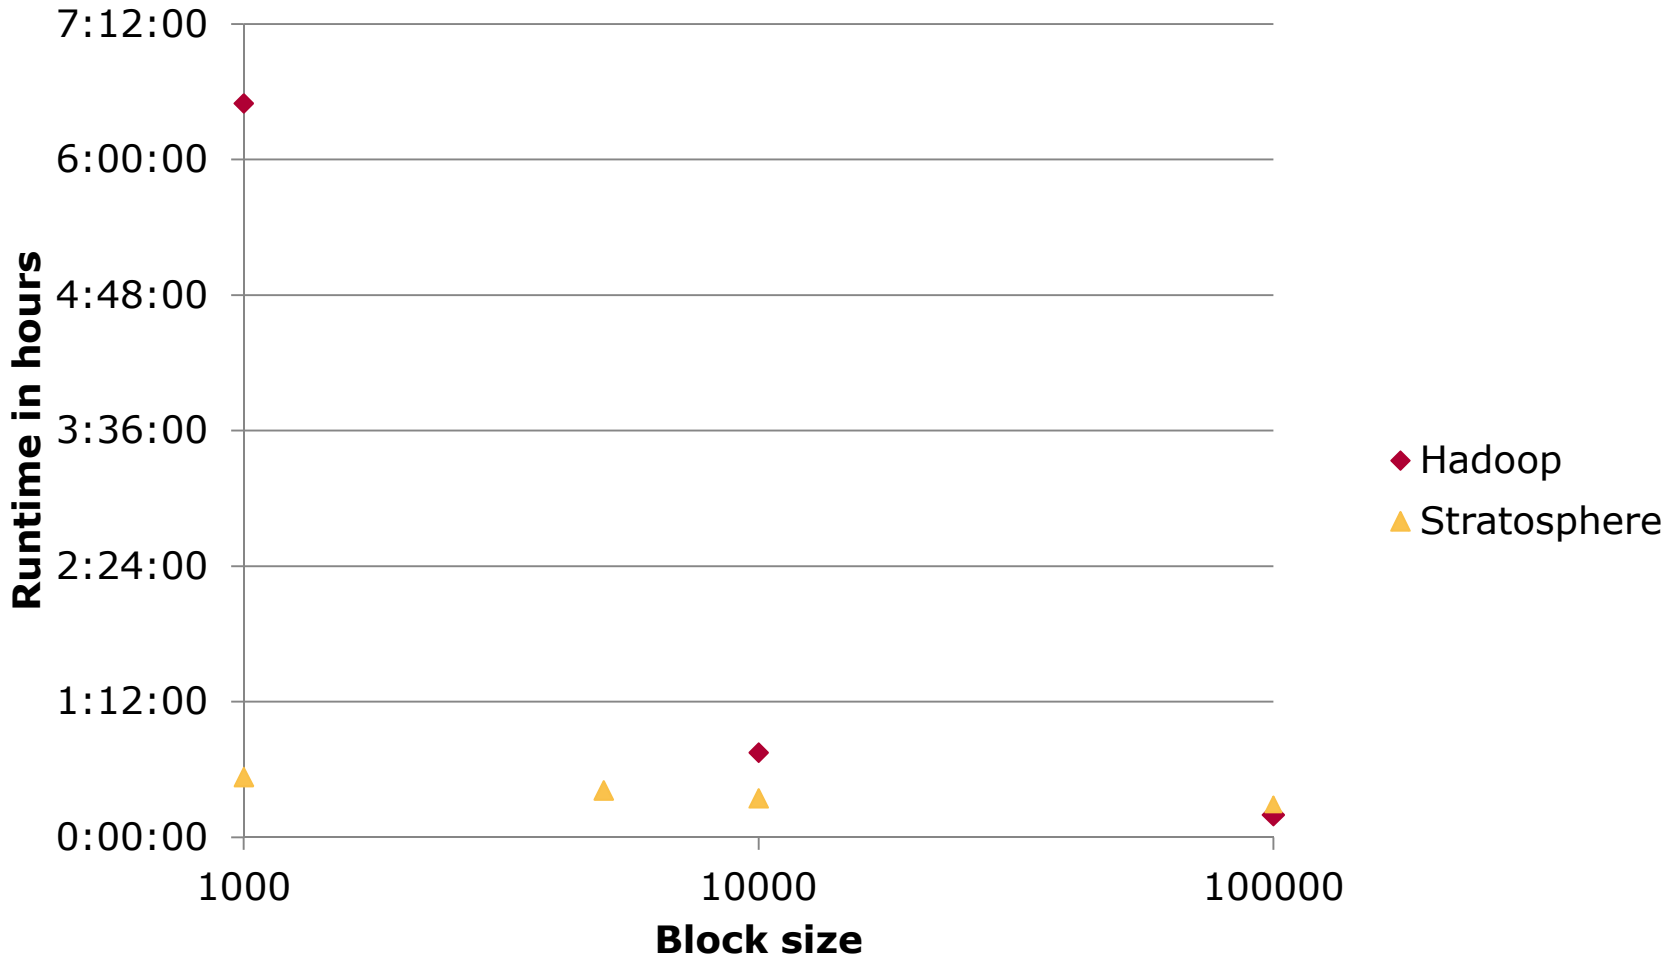

Conclusion

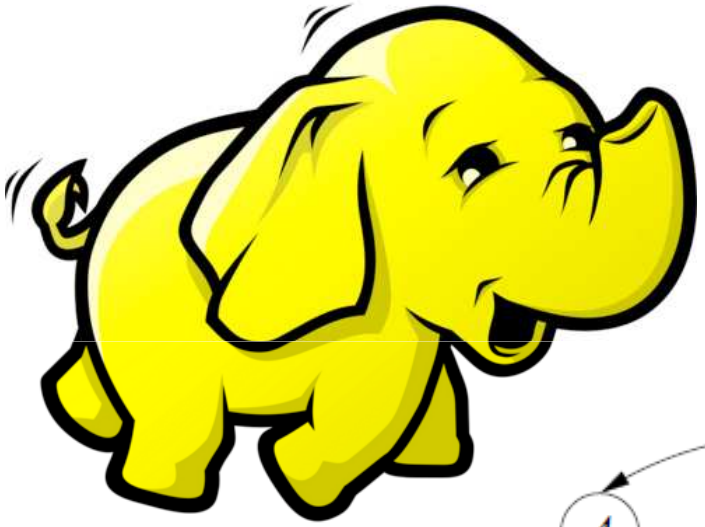

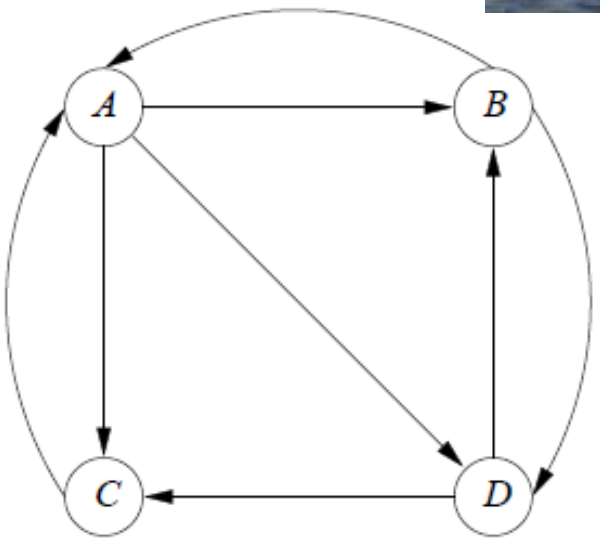

Runtime

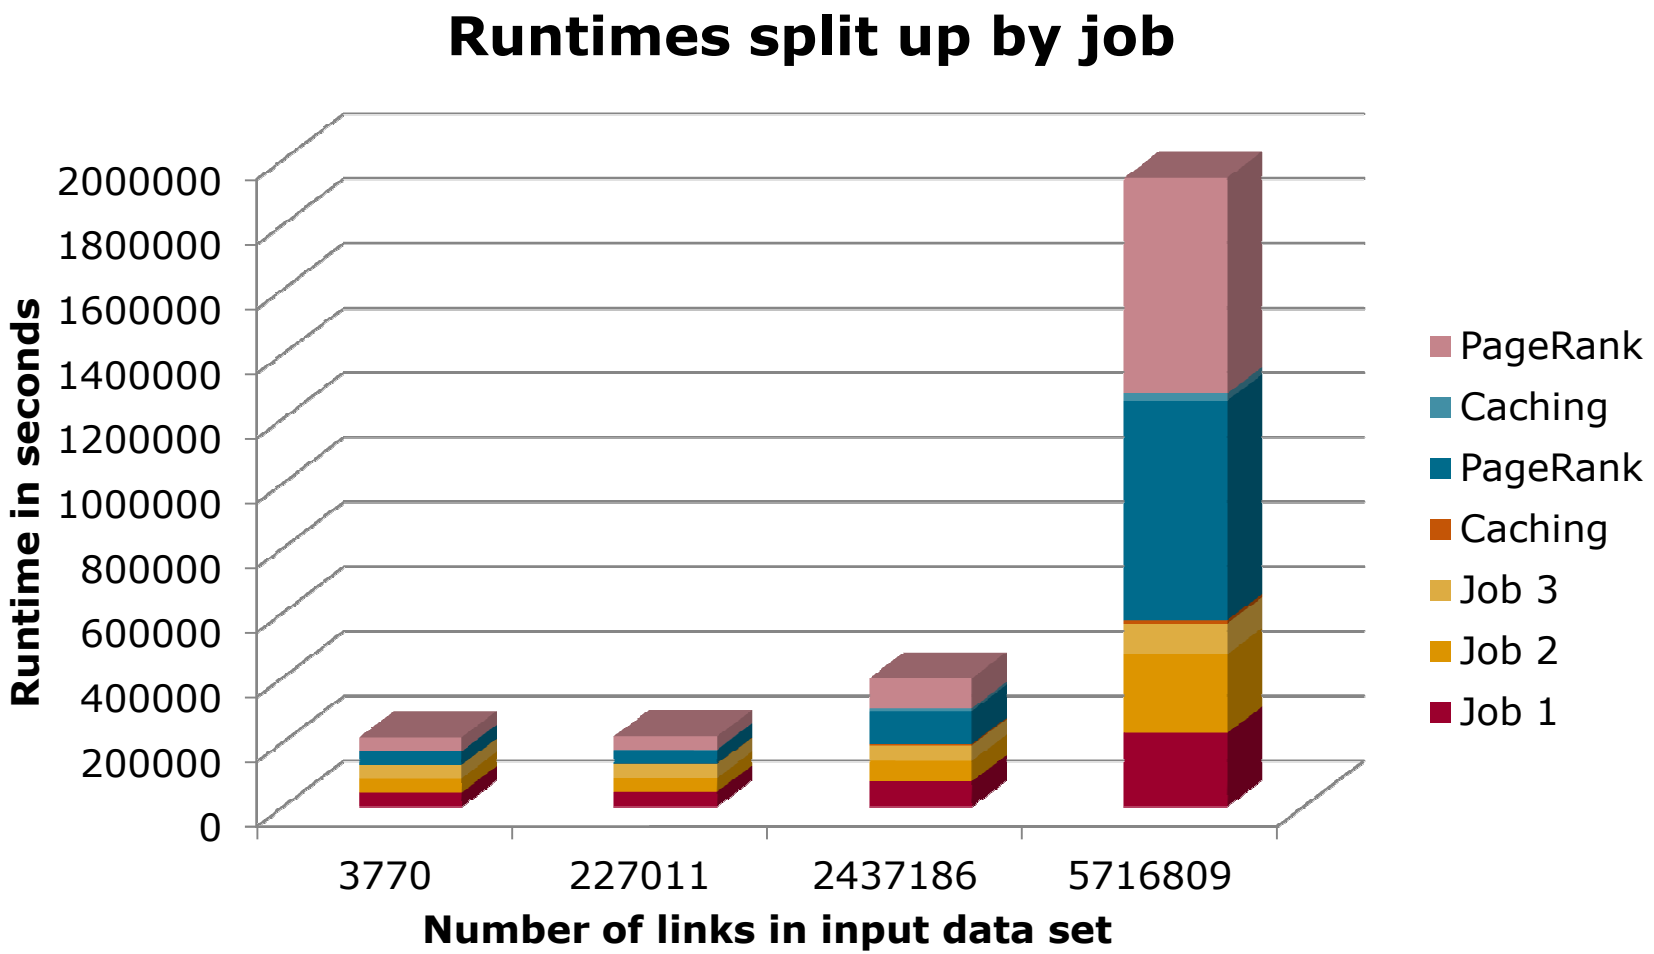
Runtimes split up by job
Runtime in seconds
2000000
1800000
1600000
1400000
1200000
1000000
800000
600000
400000
200000
0
3770
227011
2437186
5716809
Number of links in input data set
PageRank
Caching
PageRank
Caching
Job 3
Job 2
Job 1