

APACHE
LUCENE
EUROCON



Scaling search at Trovit with Solr and Hadoop

Marc Sturlese, Trovit
marc@trovit.com, 19 October 2011

Presented by

lucid
IMAGINATION



My Background



- Marc Sturlese
- Trovit
- Software engineer focused on R&D
- Responsible for search and scalability

Agenda

- What is Trovit? Why Solr and Hadoop?
- The data workflow
- Distributed indexing strategy
- Moving indexing features out of Solr
 - **Text analysis**
 - **Deduplication**
- Performance
- Questions

What is Trovit?

- Search engine for classified ads
- Tech company located in Barcelona
- Started 5 years ago just in a single country
- Now it's in 33 countries and 4 business categories
- Main purpose is to serve good quality results to the end user as fast as possible

What is Trovit?



Homes Cars Jobs Products

Homes for rent Search

Examples: [for rent terraced house leads](#), [for rent terraced house leads furnished](#)

List | Photos | Map | Favourites (0)

Login | Register

Homes for rent: terraced leads for rent

1-10 of 1,604 flats to let found

Filter your search

- Publication date
- Price
- Bedrooms
- Bathrooms
- Cities

Leeds (969)

Pudsey (25)

Bradford (14)

Wakefield (12)

Rochdale (10)

London (7)

Heywood (5)

Huddersfield (5)

Dewsbury (4)

Halifax (3)

- Regions
- Pictures
- Property features
- Property type

Continue searching

Leeds

Holiday rentals Leeds

Flatshare Leeds

Queries related to "terraced leads for rent": [for rent terraced house leads](#), [for rent terraced house leads furnished](#), [for rent terraced house leads ls6](#), [for rent terraced house leads 10](#), [for rent terraced house leads 9](#), [for rent terraced house leads ls13](#) - [see more >>](#)

[Receive new property to let by email](#)

Sort by: Relevance

Ads by Google

Flat to rent in Leeds

Quality flats and houses to rent powered by new property portal
[www.propertypicnic.com](#)

Houses to Rent

Find or List House to Rent Free Short & Long Term. Furnished & Not
[www.sublet.com](#)



House - Tong Road, Wortley, Leeds, West Yorkshire

£475 - Wortley
Two bedroom terraced house furnished, deposit? 575 gas ch, pvcu double glazing good access to leads ...
email4property

Sponsored ads

Bryan Street, Farsley, Leeds LS28, 2 bedroom terraced

Bryan Street, Farsley Farsley, Pudsey LS28 5JP | [Map](#)



available now *unfurnished* lovely two bedroom stone terrace - close to the centre of farsley village which is convenient for leads and bradford city centres. commuting... and bathroom. low maintenance paved seating area to rear of property. location farsley is a small village situated conveniently for leads and bradford...

1 week ago in zoopla! [Waite & Co](#) - [Create alert](#) - [Share](#) - [More](#)

£133

2 bedrooms
1 bathroom

Average Rent Price In Leeds



£462.27

Priced To Rent Terraced Leads



4 bedrooms
£65
↓ £60



2 bedrooms
£575
↓ £550

Why Solr and Hadoop?

- Start as custom Lucene search server
- Solr is very extensible and has a great community so we made the move
- Datastore was MySQL and custom Data Import Handler for indexing
- Scale up was not the way!
- Sharded MySQL strategies are hard to maintain
- Hadoop seemed a good feed

The data workflow

- Documents are crunched by a pipeline of MapReduce jobs
- Stats are saved for each pipeline phase to keep track of what happen every moment
- Hive is used for those stats generation

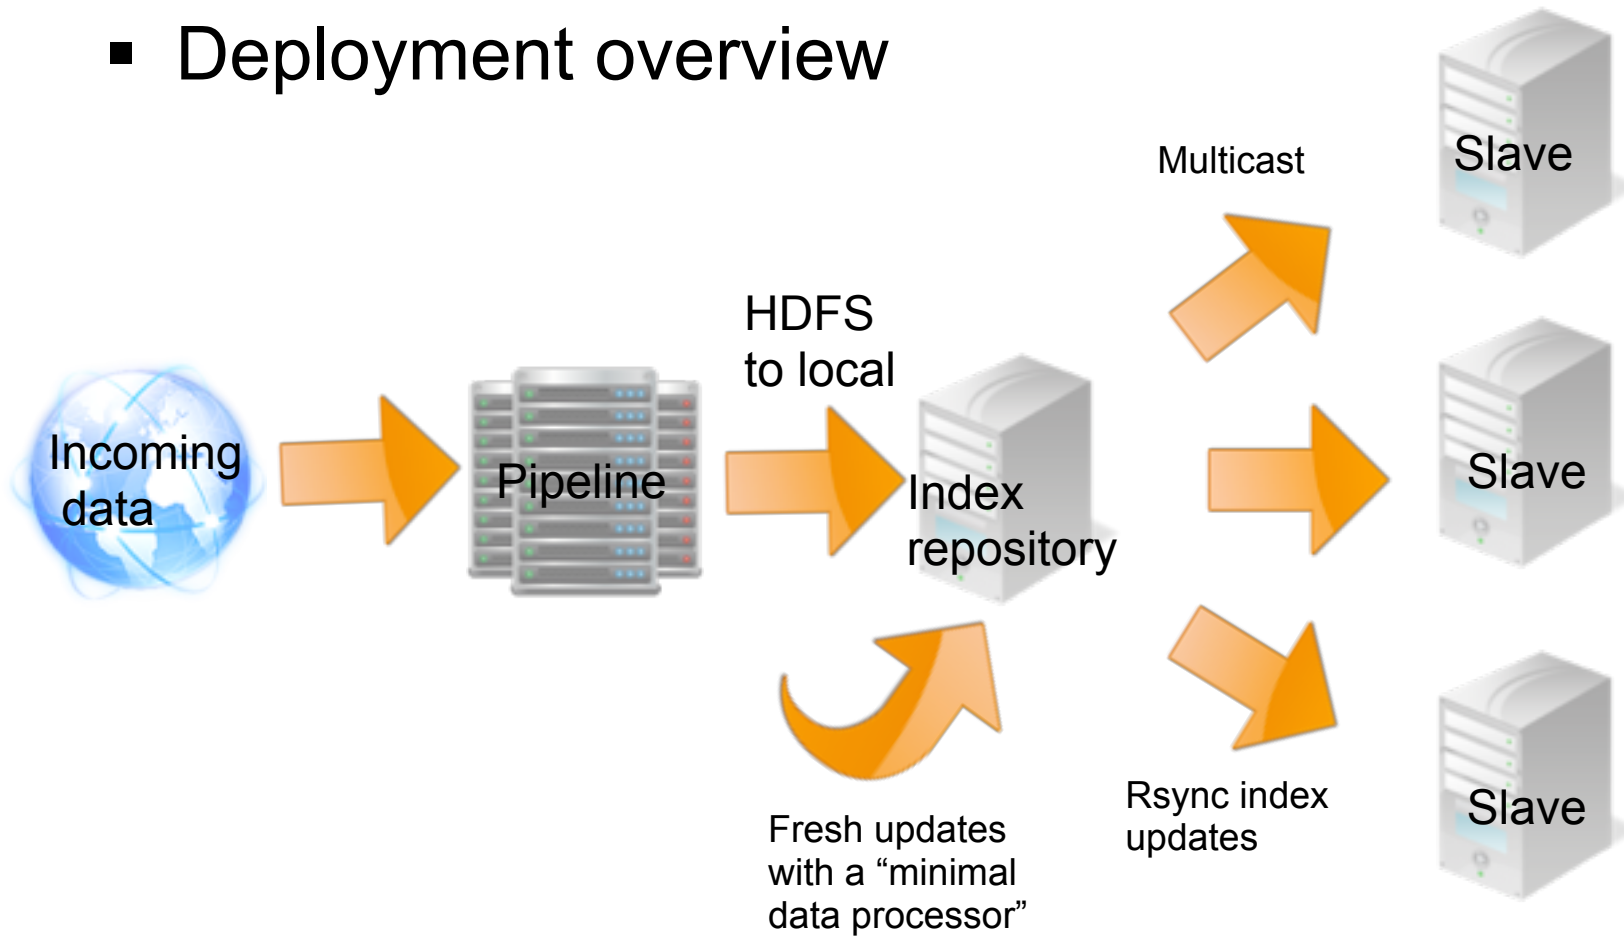
The data workflow

- Pipeline overview



The data workflow

- Deployment overview



The data workflow

- Index constantly built from scratch. Keep desired number of segments for good search performance
- “Minimal data processor” allows fresh data appear in the search results
- HDFS makes backups really convenient
- Multicast system allows to send indexes to all slaves at the same time. The only limit is your bandwidth

Distributed indexing strategy

- First looked at SOLR-1301
- Extends InputFormat allowing just an index per reducer
- Good to generate huge indexes building a shard per reduce
- To achieve the goal with minimal time, shards should have very similar size
- Reduce side indexing seemed the way but... indexes differ a lot of size depending on the country and vertical

Distributed indexing strategy

- Single monolithic indexes or shards
- Another approach, 2 phases indexing (2 sequential MapReduce jobs)
 - **Partial indexing: Generate lots of “micro indexes” per each monolithic or sharded index**
 - **Merge: Groups all the “micro indexes” and merge them to get the production data.**

Distributed indexing strategy

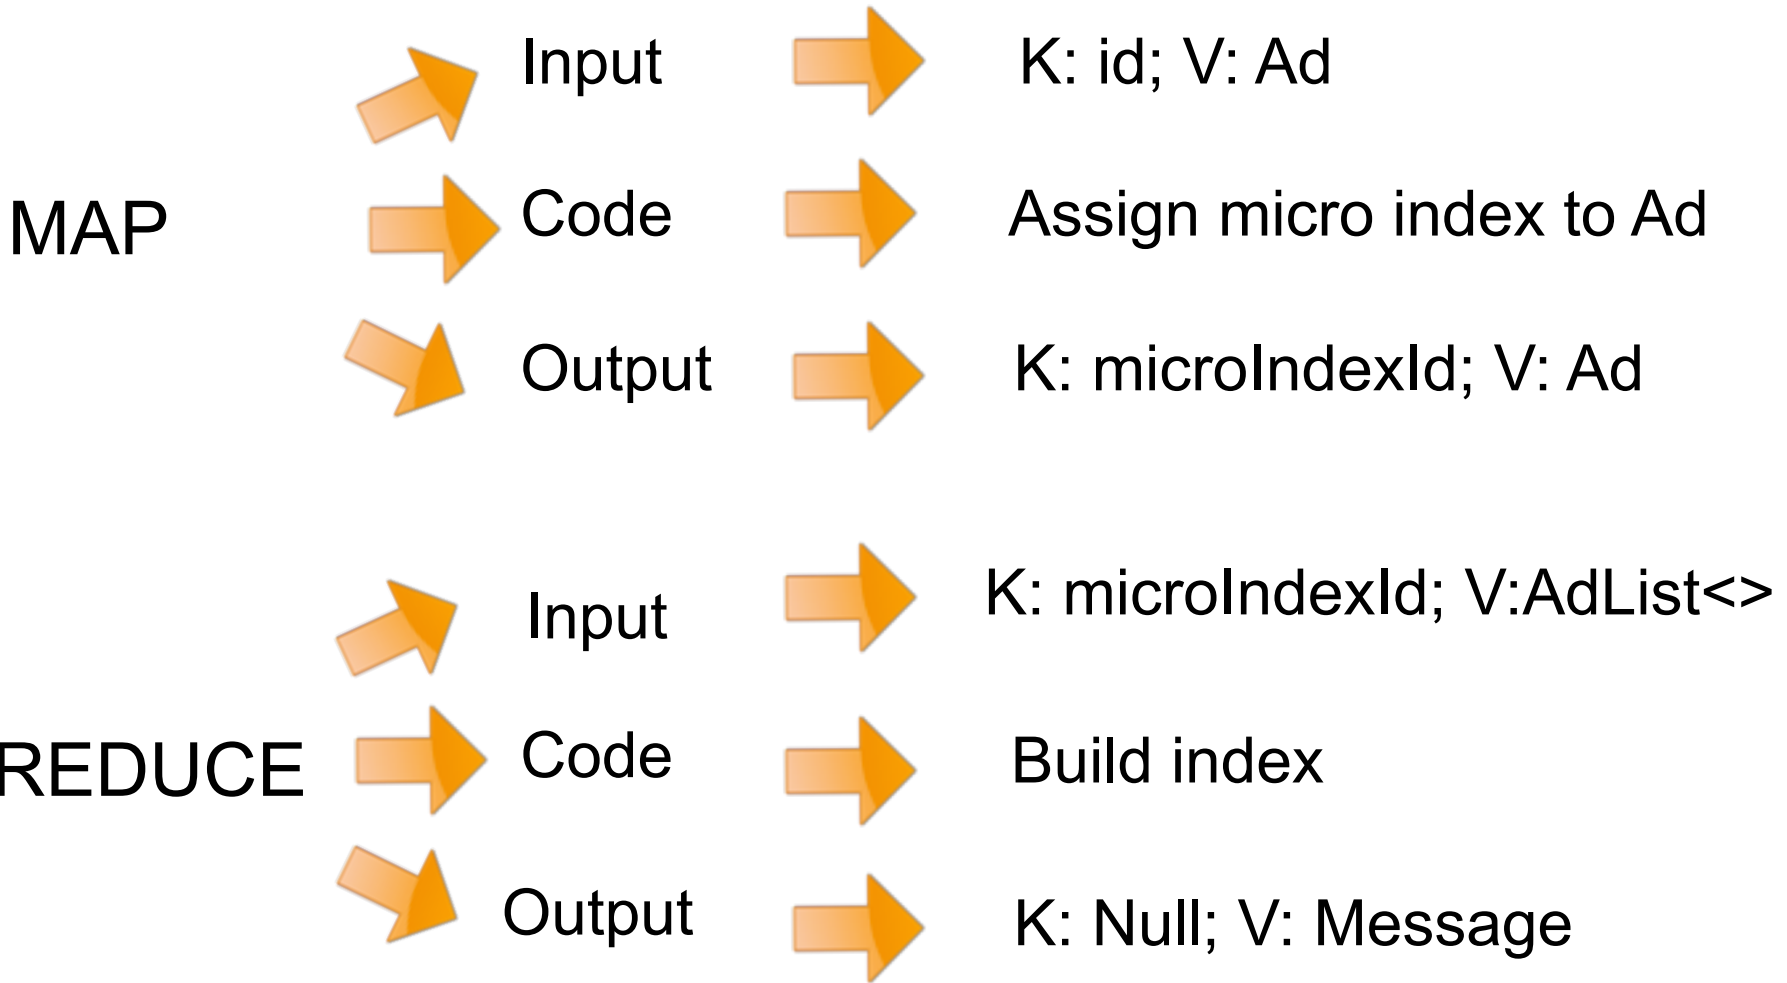
- 2 phases indexing overview



Distributed indexing - Partial generation

- Map reads serialized data and emit grouping by micro index
- Reduce receives ads grouped as “micro index” and builds it
- Embedded Solr Server for indexing and optimize
- Solr cores configuration is stored in HDFS.
- Indexes are build in local disk and then uploaded to HDFS

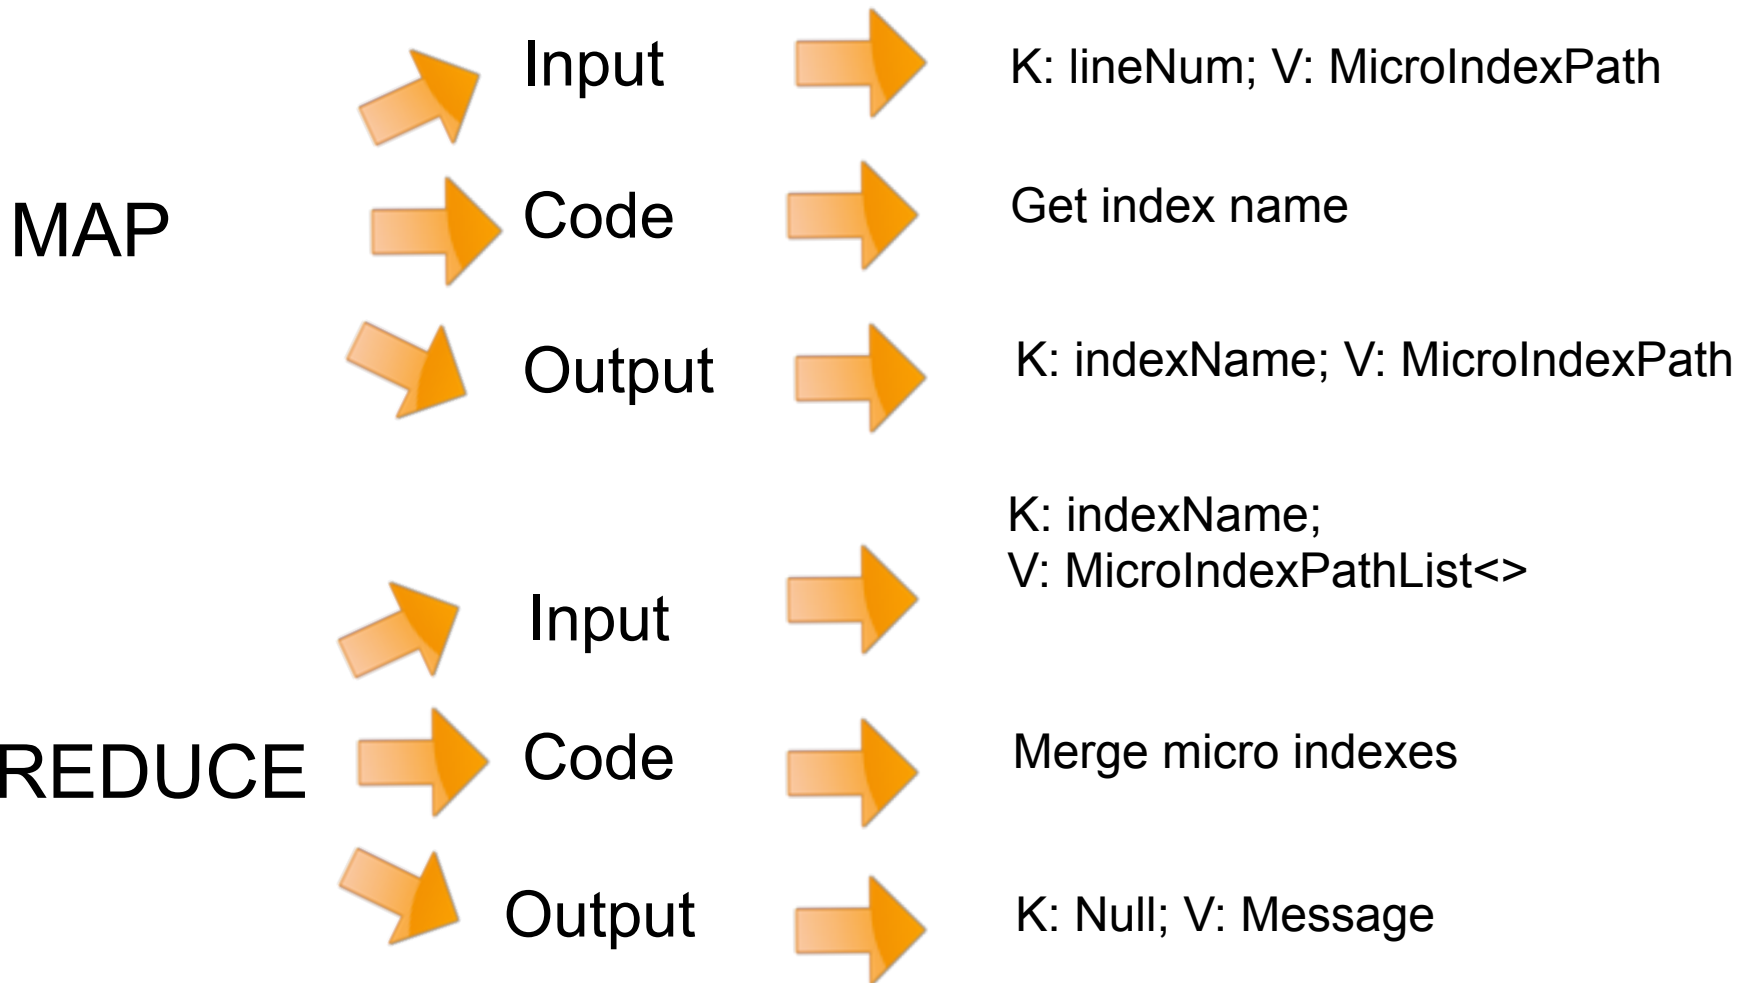
Distributed indexing - Partial generation



Distributed indexing - Merge phase

- Merging is done in plain Lucene
- Map reads a list of the paths of the micro indexes and emits grouping per shard or monolithic
- Reduce receives the list and does the proper merge
- Partial indexes are downloaded to local disk, merged to a single one and uploaded back to HDFS
- Since Lucene 3.1 `addIndexes(Directory)` uses copy, merge can be very fast

Distributed indexing – Merge phase



Distributed indexing strategy

- Pros:
 - Highly scalable
 - Allows indexes with very different size keeping good performance
 - Easy to manage
- Cons:
 - Time uploading and downloading from HDFS before it gets into production

Moving features out of Solr

- Useful when you have to deal with lots of data
- Text processing with Solr and Hadoop
- Distributing Solr Deduplication

Text processing with Solr and Hadoop

- Solr has many powerful analyzers already implemented
- Mahout tokenizes text using plain Lucene and Hadoop
- The setUp method on a Map instantiates the Analyzer
- A Map receives serialized data and that is processed using Solr analyzers
- Analyzer can receive configuration parameters from a job-site.xml file

Text processing with Solr and Hadoop



```
//init Solr analyzer

final List<TokenFilterFactory> filters = new ArrayList<TokenFilterFactory>();

    final TokenFilterFactory wordDelimiter = new WordDelimiterFilterFactory();

    Map<String, String> args = new HashMap<String, String>();

    args.put("generateWordParts", conf.get(WORD_PARTS));

    args.put("splitOnNumerics", conf.get(NUMERIC_SPLIT));

    wordDelimiter.init(args);

    final TokenFilterFactory accent = new ISOLatin1AccentFilterFactory();

    final TokenFilterFactory lowerCase = new LowerCaseFilterFactory();

    filters.add(wordDelimiter);

    filters.add(accent);

    filters.add(lowerCase);

    final TokenizerFactory tokenizer = new StandardTokenizerFactory();

    analyzer = new TokenizerChain(null, tokenizer, filters.toArray(new
    TokenFilterFactory[filters.size()]));
```

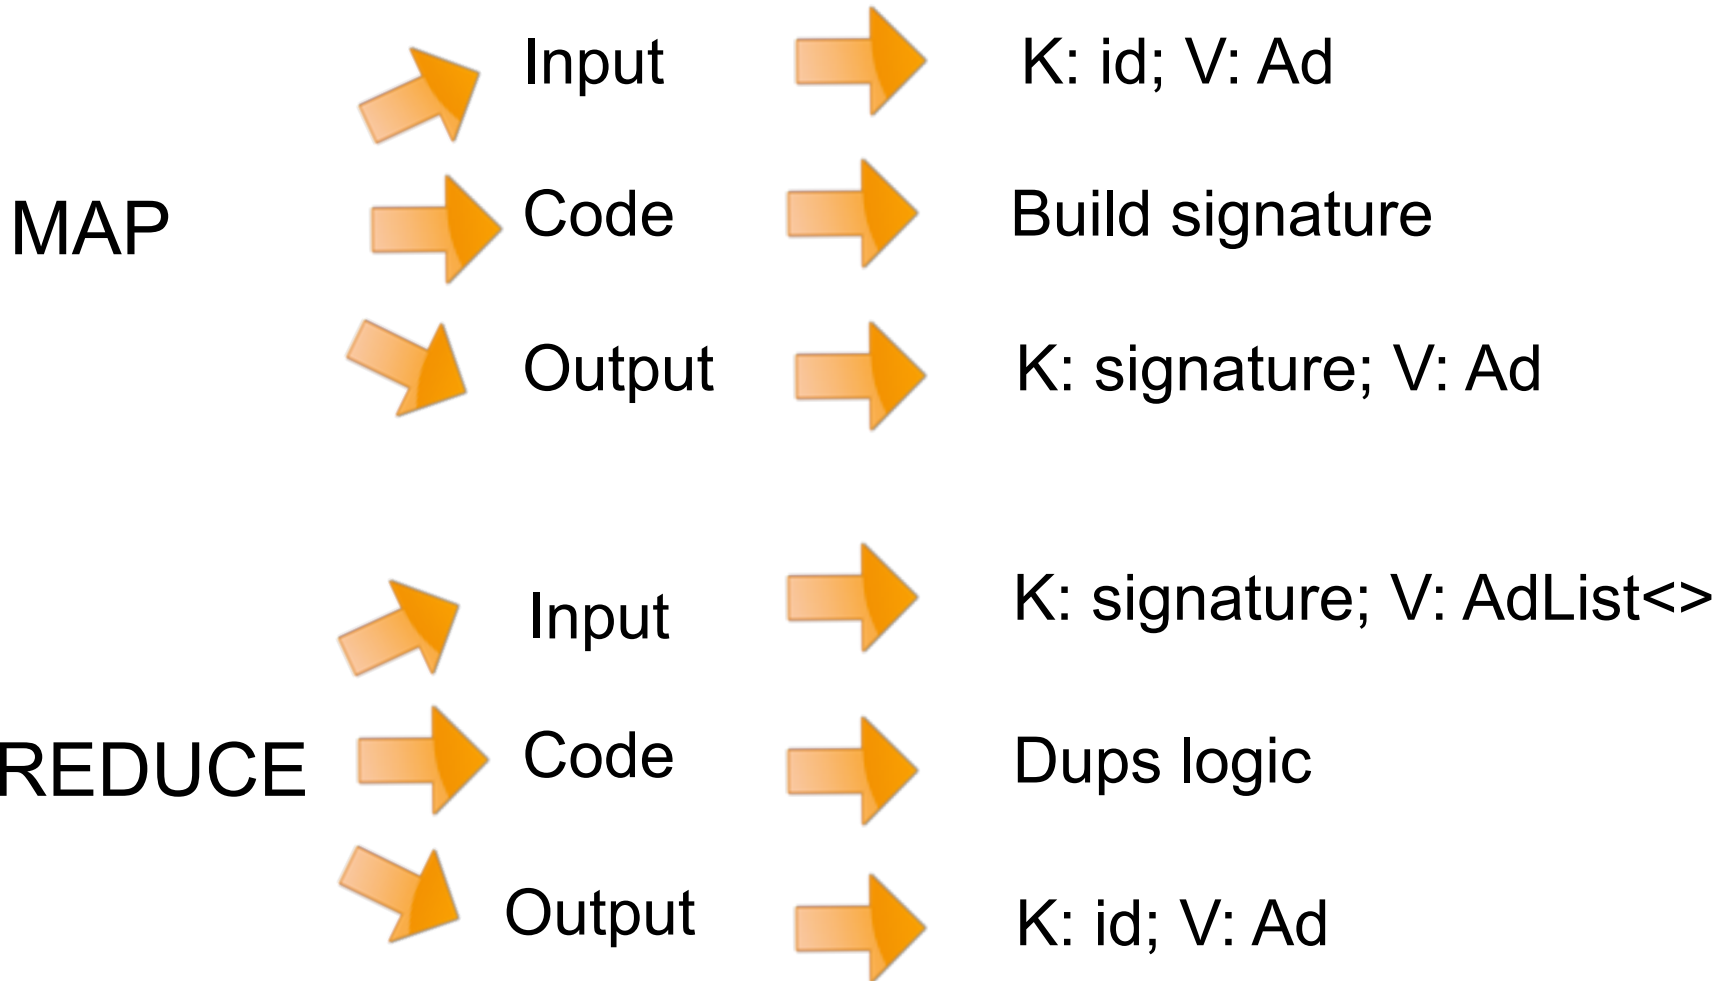
Text processing with Solr and Hadoop

```
//Tokenizing text  
  
...  
  
HashSet<String> tokens = new HashSet<String>();  
  
TokenStream stream = analyzer.reusableTokenStream(fieldName, new  
StringReader(fieldValue));  
  
TermAttribute termAtt = (TermAttribute) stream.addAttribute(TermAttribute.class);  
  
while (stream.incrementToken()) {  
    tokens.add(termAtt.term());  
}  
  
return tokens;
```

Distributed deduplication

- Compute near duplicates in a distributed environment
- Map receives serialized ads and emit building the key using Solr's TextProfileSignature
- Reduce receives dup ads grouped. There, you decide
- Field names to compute the signature received as configuration parameters from a job-site.xml file

Distributed deduplication



Performance: Setting the merge factor

- Used in LogByteMergePolicy and older
- When indexing, tells Lucene how many segments can be created before a merge happen
- Very low value will keep the index almost optimized. Good for search performance but indexing will be slower
- High value will generate lots, of files. Indexing will be faster but not search requests
- New versions of Solr default to TieredMergePolicy which don't use it



Contact



- Thank you for your attention
- Marc Sturlese
 - marc@trovit.com
 - www.trovit.com