# Map-Reduce

John Hughes

# The Problem

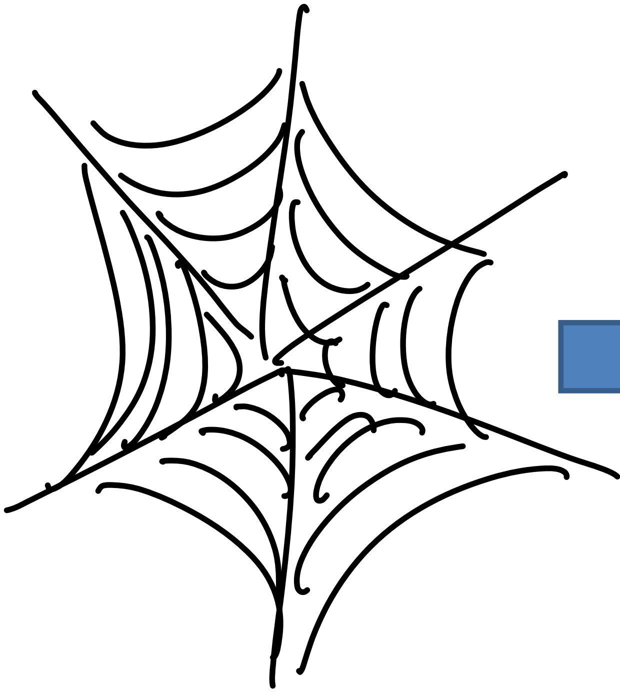

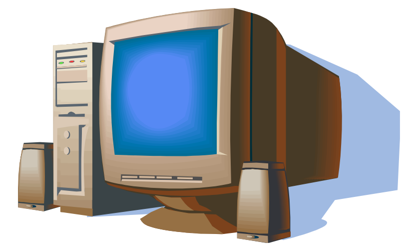

850TB
in 2006

# The Solution?

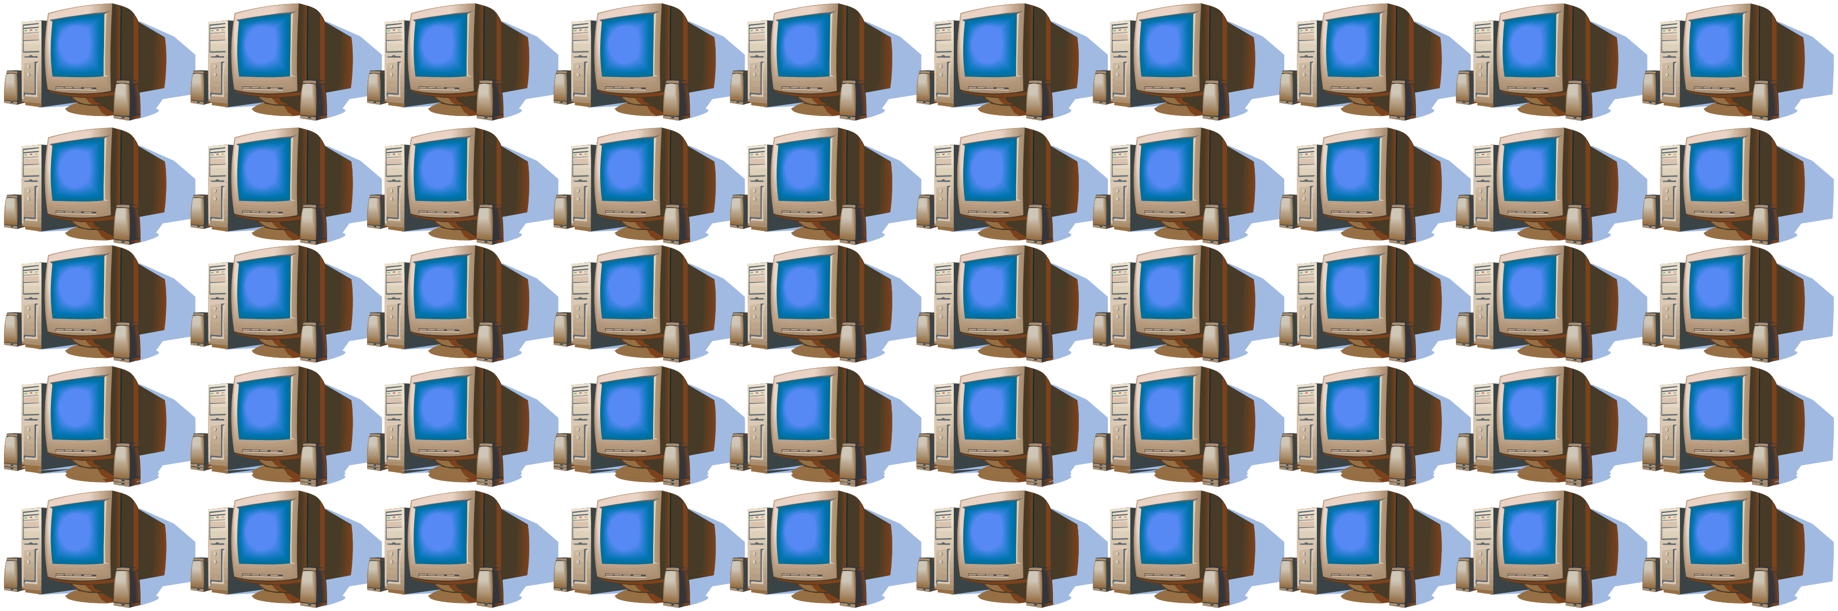

- Thousands of commodity computers networked together
- 1,000 computers → 850GB each
- How to make them work together?

# Early Days

- Hundreds of ad-hoc distributed algorithms
  - Complicated, hard to write
  - Must cope with fault-tolerance, load distribution, …

**MapReduce: Simplified Data Processing on Large Clusters**

*by Jeffrey Dean and Sanjay Ghemawat*

In Symposium on Operating Systems Design & Implementation (OSDI 2004)

# The Idea

- Many algorithms apply the *same* operation to a lot of data items, then *combine* results

- Cf map :: (a->b) -> [a] -> [b]
- Cf foldr :: (a->b->b) -> b -> [a] -> b
  - Called *reduce* in LISP

- Define a *higher-order function* to take care of distribution; let users just write the functions passed to map and reduce

# Pure functions are great!

- They can be *run anywhere* with the same result—easy to distribute
- They can be *reexecuted* on the same data to recreate results lost by crashes

# ”It’s map and reduce, but not as we know them Captain”

- Google map and reduce work on collections of *key-value pairs*

- map_reduce *mapper reducer* :: [(k,v)] -> [(k2,v2)]
  - *mapper* :: k -> v -> [(k2,v2)]
  - *reducer* :: k2 -> [v2] -> [(k2,v2)]

All the values with the same key are collected

Usually just 0 or 1

# Example: counting words

- Input: (file name, file contents)

mapper

- Intermediate pairs: (word, 1)

reducer

- Final pairs: (word, total count)

# Example: counting words

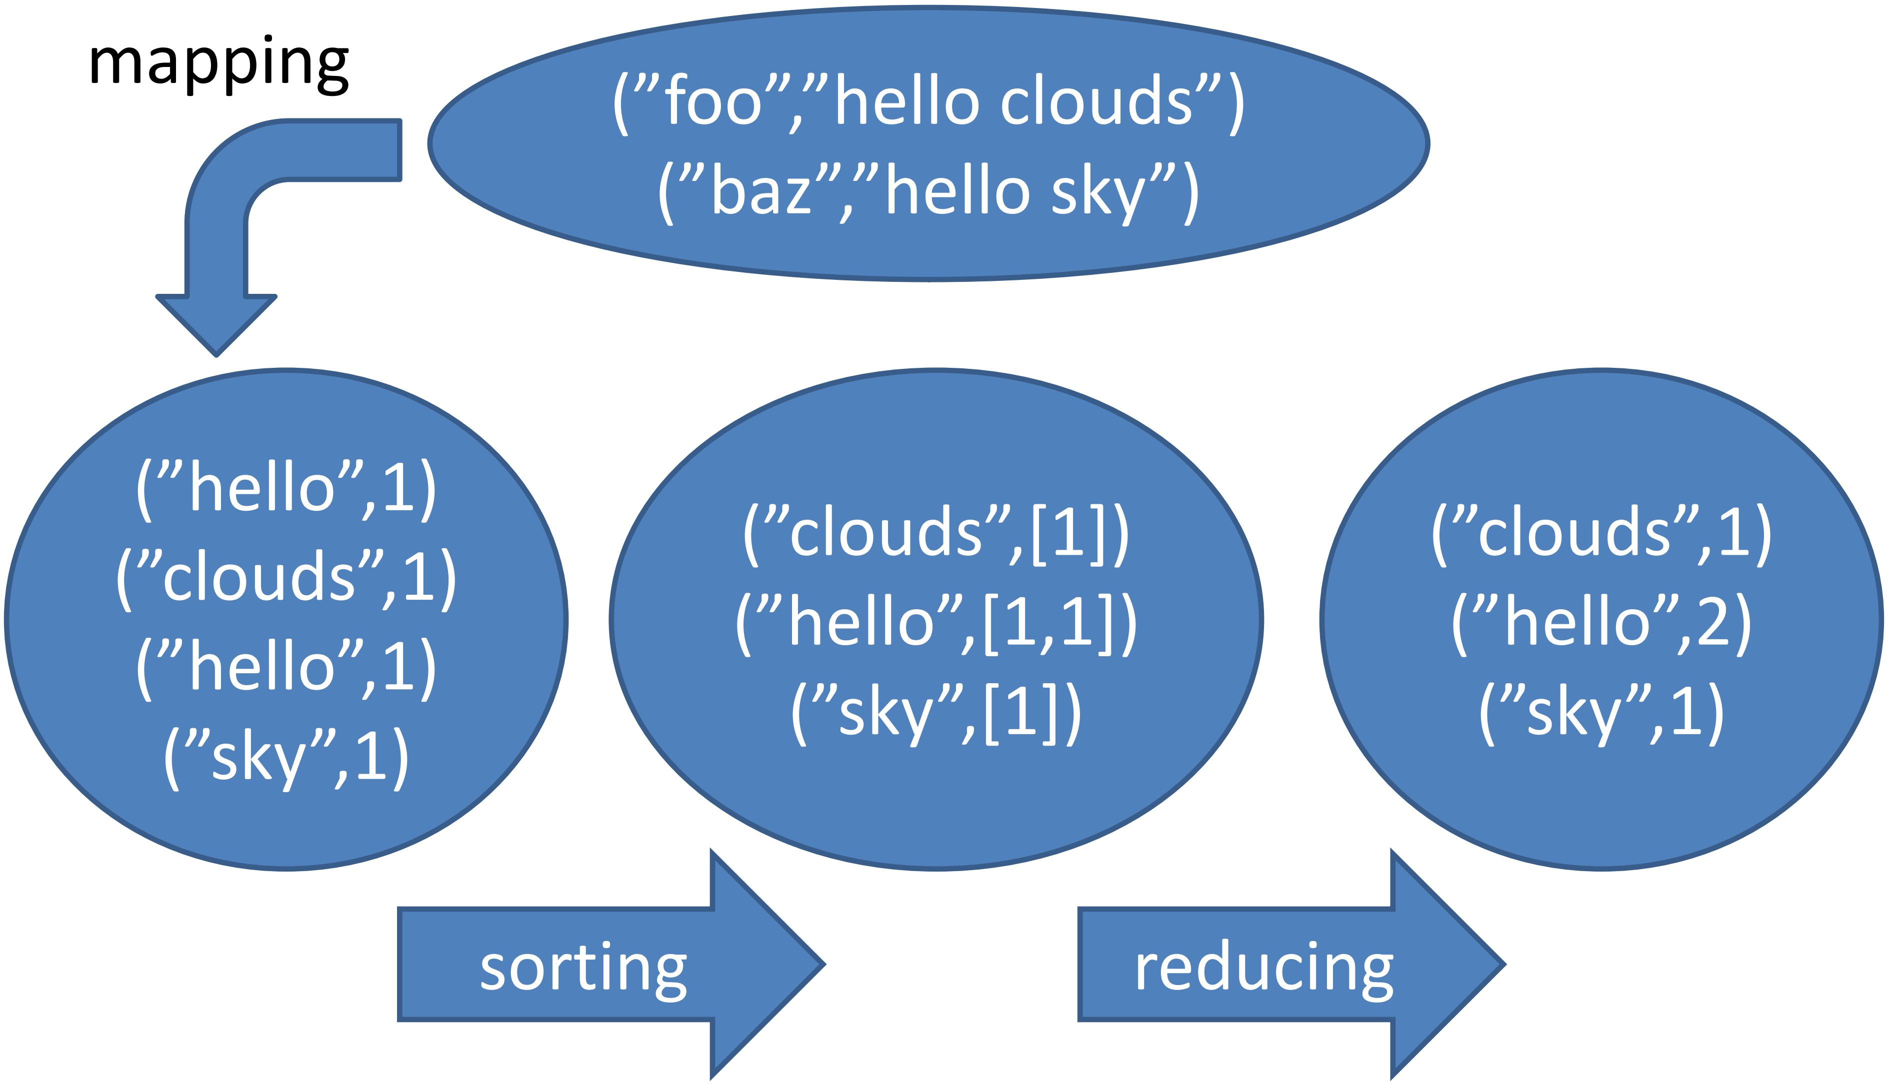

# Parallelising Map-Reduce

- Divide the input into M chunks, map in parallel
  - About 64MB per chunk is good!
  - Typically M ~ 200,000 on 2,000 machines (~13TB)

- Divide the intermediate pairs into R chunks, reduce in parallel
  - Typically R ~ 5,000

**Problem:** all {K,V} with the same key must end up in the same chunk!

# Chunking Reduce

- All pairs with the same key must end up in the same chunk

- *Map keys to chunk number: 0..R-1*
  - e.g. hash(Key) rem R

erlang:phash2(Key,R)

- Every mapper process generates inputs for all R reducer processes

# Usage

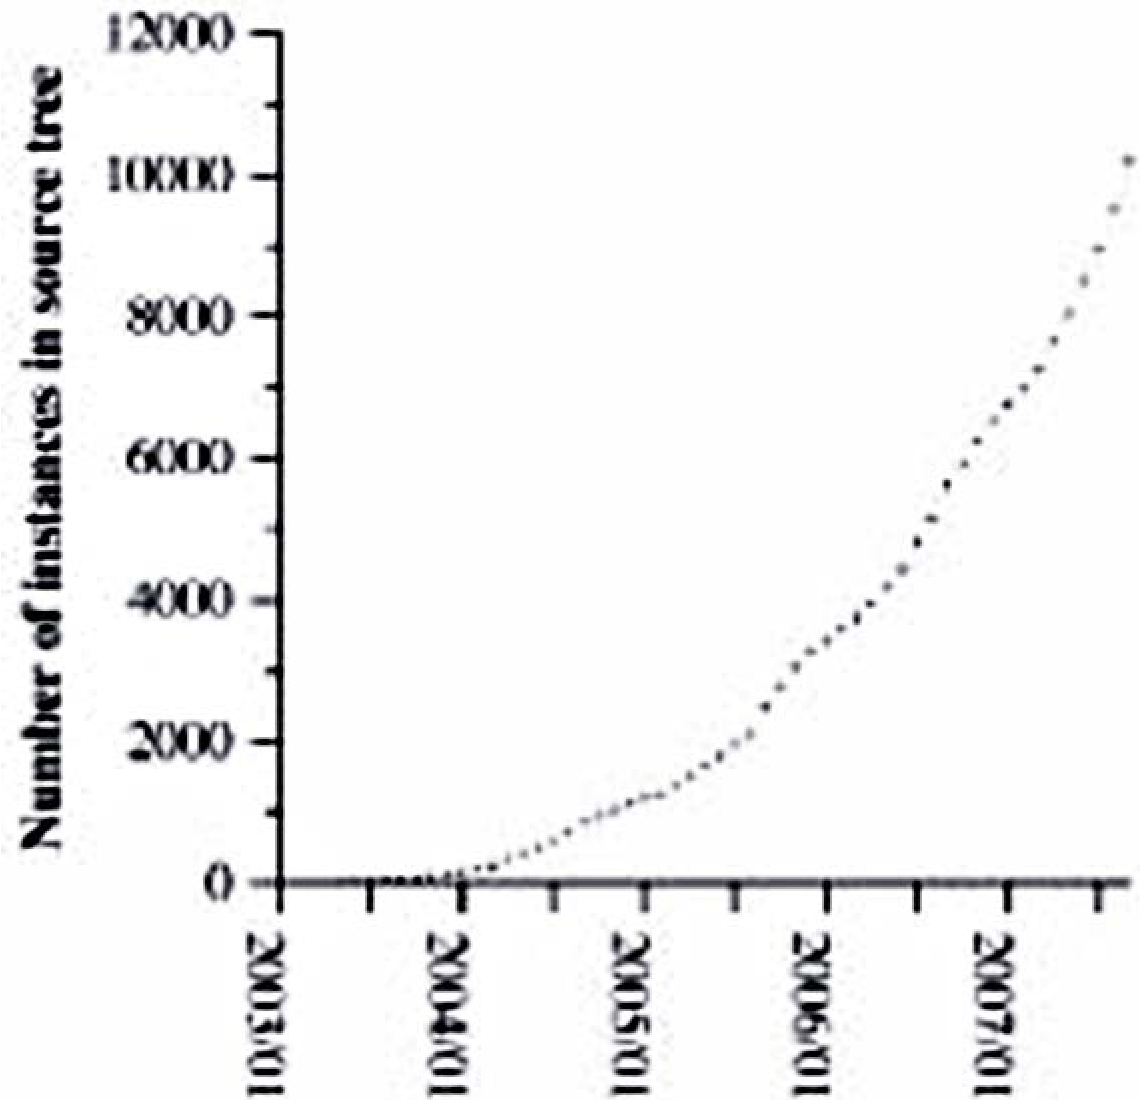

Number of instances in source tree
12000
10000
8000
6000
4000
2000
0
2003/01
2004/01
2005/01
2006/01
2007/01

# Google web search indexing

**Before**

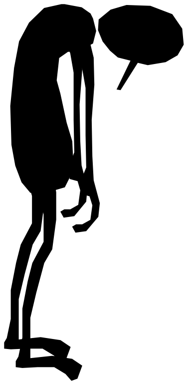

3800

LOC

**After**

700

LOC

# Experience

“Programmers find the system easy to use: more than ten thousand distinct MapReduce programs have been implemented internally at Google over the past four years, and an average of one hundred thousand MapReduce jobs are executed on Google’s clusters every day, processing a total of more than twenty petabytes of data per day.”

*From MapReduce: Simplified Data Processing on Large Clusters*
*by Jeffrey Dean and Sanjay Ghemawat, CACM 2008*

# Applications

- large-scale machine learning
- clustering for Google News and Froogle
- extracting data to produce reports of popular queries
  - e.g. Google Zeitgeist and Google Trends
- processing of satellite imagery
- language model processing for statistical machine translation
- large-scale graph computations.
- Apache Hadoop

# You may have seen…

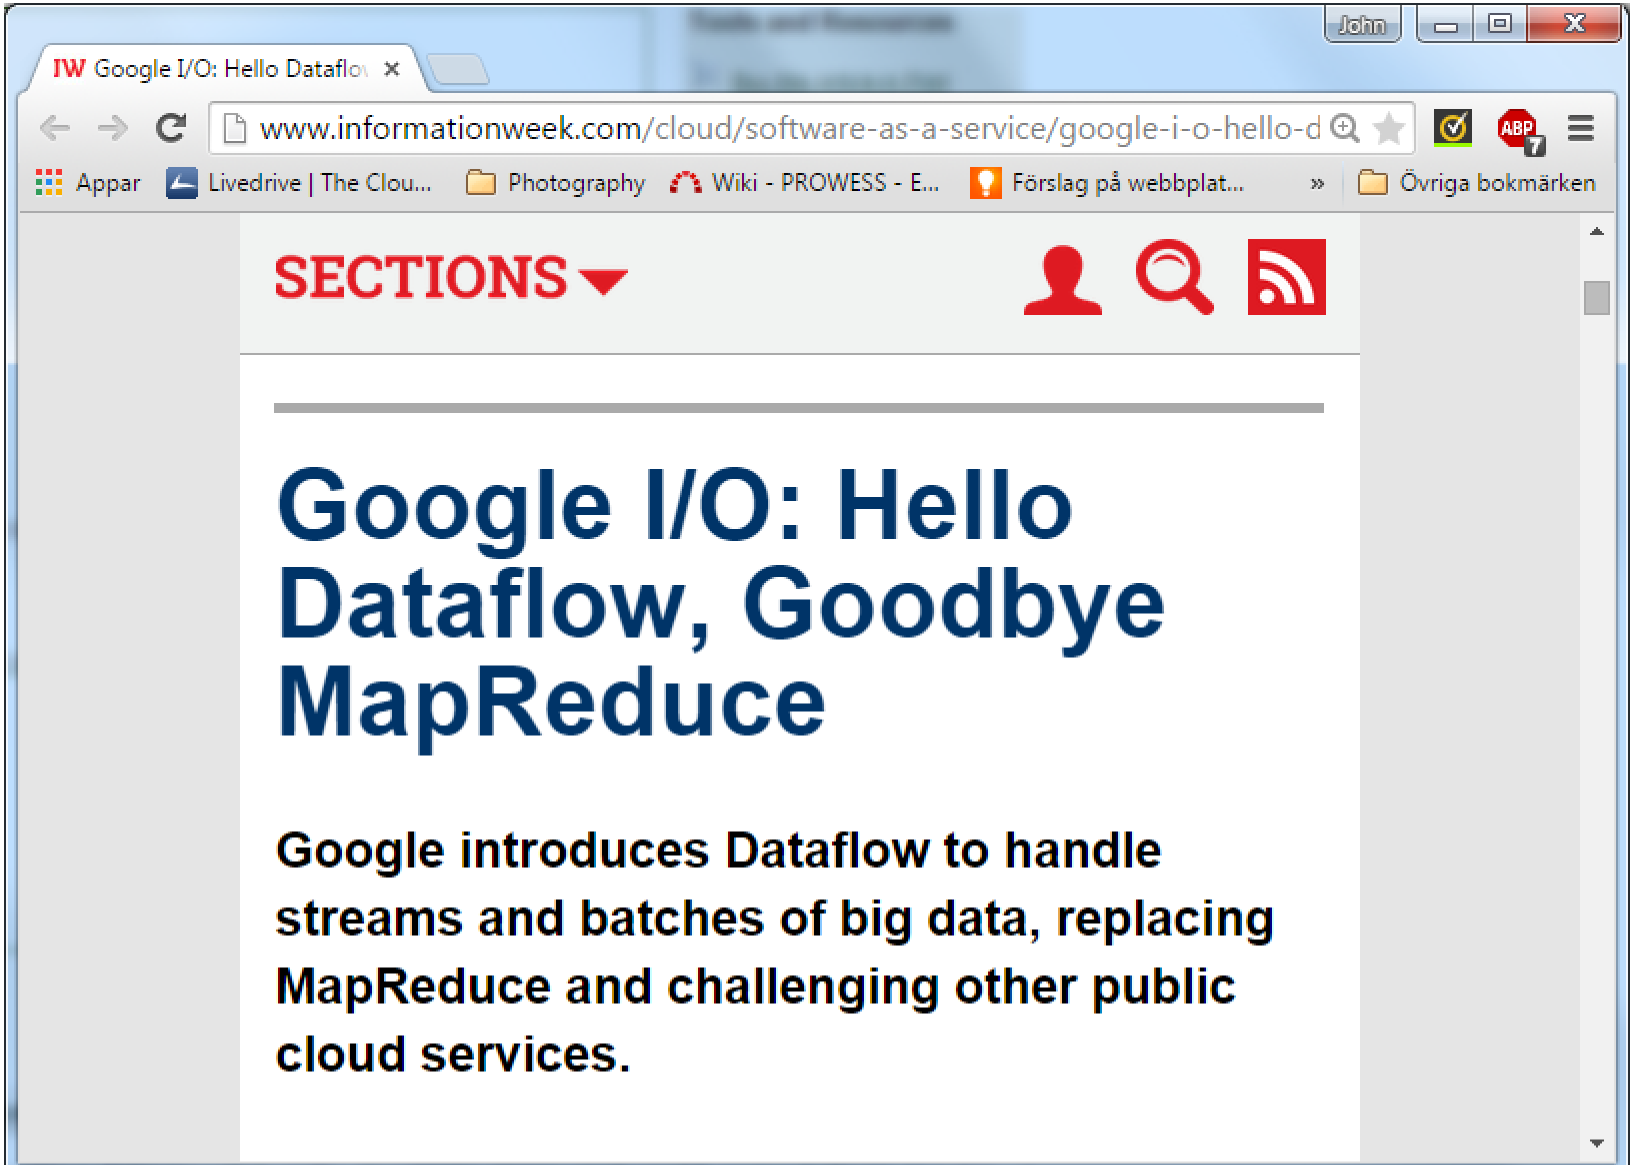

# What is it?

PLDI 2010

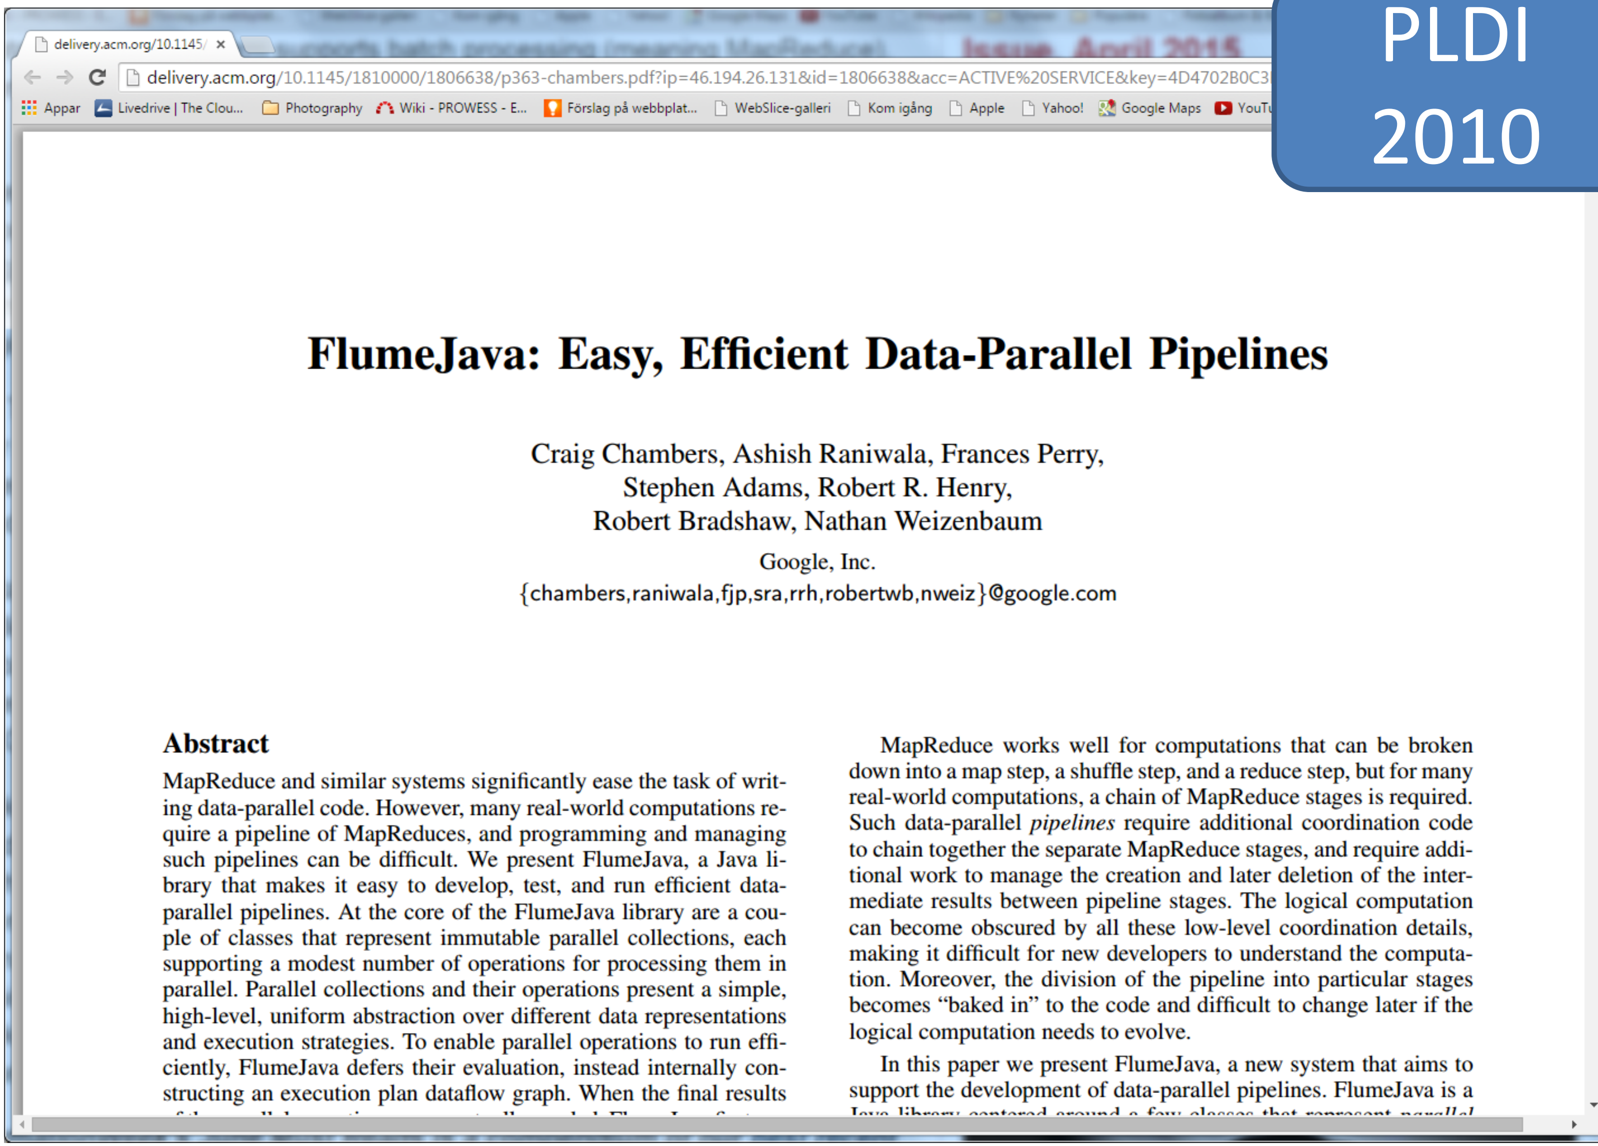

# FlumeJava: Easy, Efficient Data-Parallel Pipelines

Craig Chambers, Ashish Raniwala, Frances Perry,
Stephen Adams, Robert R. Henry,
Robert Bradshaw, Nathan Weizenbaum

Google, Inc.
{chambers,raniwala,fjp,sra,rrh,robertwb,nweiz}@google.com

## Abstract

MapReduce and similar systems significantly ease the task of writing data-parallel code. However, many real-world computations require a pipeline of MapReduces, and programming and managing such pipelines can be difficult. We present FlumeJava, a Java library that makes it easy to develop, test, and run efficient data-parallel pipelines. At the core of the FlumeJava library are a couple of classes that represent immutable parallel collections, each supporting a modest number of operations for processing them in parallel. Parallel collections and their operations present a simple, high-level, uniform abstraction over different data representations and execution strategies. To enable parallel operations to run efficiently, FlumeJava defers their evaluation, instead internally constructing an execution plan dataflow graph. When the final results

MapReduce works well for computations that can be broken down into a map step, a shuffle step, and a reduce step, but for many real-world computations, a chain of MapReduce stages is required. Such data-parallel *pipelines* require additional coordination code to chain together the separate MapReduce stages, and require additional work to manage the creation and later deletion of the intermediate results between pipeline stages. The logical computation can become obscured by all these low-level coordination details, making it difficult for new developers to understand the computation. Moreover, the division of the pipeline into particular stages becomes "baked in" to the code and difficult to change later if the logical computation needs to evolve.

In this paper we present FlumeJava, a new system that aims to support the development of data-parallel pipelines. FlumeJava is a

# What is it?

- A datatype of *immutable parallel collections*
  - which can be distributed over a data centre
  - or consist of *streaming data*
- An API including *map*, *reduce, filter*, *group*… that apply *pure functions* to collections
- An *optimising on-the-fly compiler* that converts FlumeJava pipelines to a sequence of MapReduce jobs…
- A *higher-level interface* built on top of MapReduce