

Computing Recommendations at Extreme Scale with Apache Flink



Till Rohrmann

Flink committer / data Artisans

trohrmann@apache.org

[@stsffap](https://twitter.com/stsffap)

Recommendations: Collaborative Filtering

Recommendations



- Omnipresent nowadays
- Important for user experience and sales

amazon.com

Recommended for You


Amazon.com has new recommendations for you based on [items](#) you purchased or told us you own.



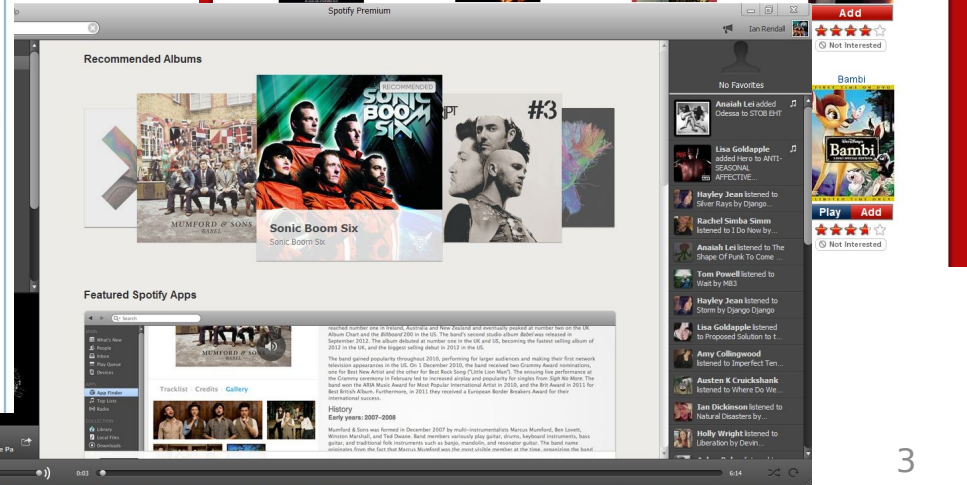
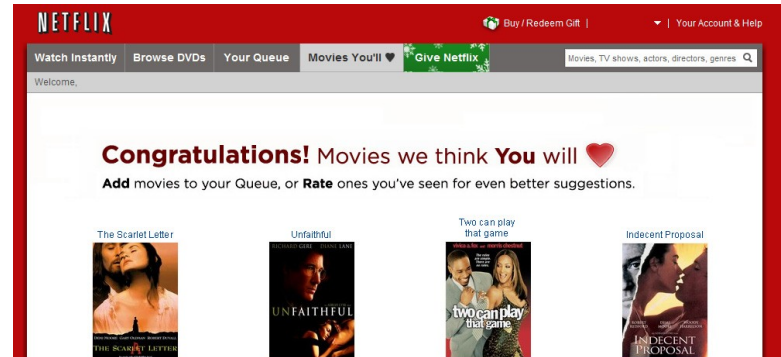
[Google Apps Deciphered: Compute in the Cloud to Streamline Your Desktop](#)




[Google Apps Administrator Guide: A Private-Label Web Workspace](#)



[Googlepedia: The Ultimate Google Resource \(3rd Edition\)](#)



Recommending Websites



Collaborative Filtering



- Recommend items based on users with similar preferences
- Latent factor models capture underlying characteristics of items and preferences of user
- Predicted preference: $\hat{r}_{u,i} = \mathbf{x}_u^T \mathbf{y}_i$

Rating Matrix



- Explicit or implicit ratings
- Prediction goal: Rating for unseen items

Items

<i>Users</i>	10	5	?	⋮
	?	2	10	⋮
	7	?	5	⋮
		Facebook		Princess

Matrix Factorization



- Calculate low rank approximation to obtain latent factors

$$\begin{matrix} & \text{Items} \\ \text{Users} & \begin{pmatrix} 10 & 5 & ? \\ ? & 2 & 10 \\ 7 & ? & 5 \end{pmatrix} \approx \begin{pmatrix} X \\ X \\ X \end{pmatrix} \bullet \begin{pmatrix} Y \\ Y \\ Y \end{pmatrix}
 \end{matrix}$$

Facebook

Princess

$$\min_{X, Y} \sum_{r_{u,i} \neq 0} (r_{u,i} - x_u^T y_i)^2 + \lambda \left(\sum_u n_u \|x_u\|^2 + \sum_i n_i \|y_i\|^2 \right)$$

- $r_{u,i}$ rating of user u for item i
- x_u latent factors of user u
- y_i latent factors of item i
- λ regularization constant
- n_u number of rated items for user u
- n_i number of ratings for item i

Alternating Least Squares



- Hard to optimize since we have two variables X, Y
- Fixing one variables gives quadratic problem

$$x_u = (YS^u Y^T + \lambda n_u I)^{-1} Y r_u^T$$

$$S_{ii}^u = \begin{cases} 1 & \text{if } r_{ui} \neq 0 \\ 0 & \text{else} \end{cases}$$

We only need the item vectors rated by user u

ALS Algorithm



- Update user matrix

Items Calculate update

$$\begin{matrix} \text{Users} & \begin{pmatrix} 10 & 5 & ? \\ ? & 2 & 10 \\ 7 & ? & 5 \end{pmatrix} & \approx & \begin{pmatrix} X \\ X \\ X \end{pmatrix} & \bullet & \begin{pmatrix} Y \\ Y \\ Y \end{pmatrix} \end{matrix}$$

Keep fixed

The diagram illustrates the ALS update step for the user matrix. It shows a 3x3 user-item rating matrix with some missing values (indicated by '?'). This matrix is approximated by the product of a user matrix X (a 3x1 column vector) and an item matrix Y (a 1x3 row vector). The user matrix X is highlighted with a green box and labeled "Calculate update", while the item matrix Y is highlighted with a red box and labeled "Keep fixed".

ALS Algorithm contd.



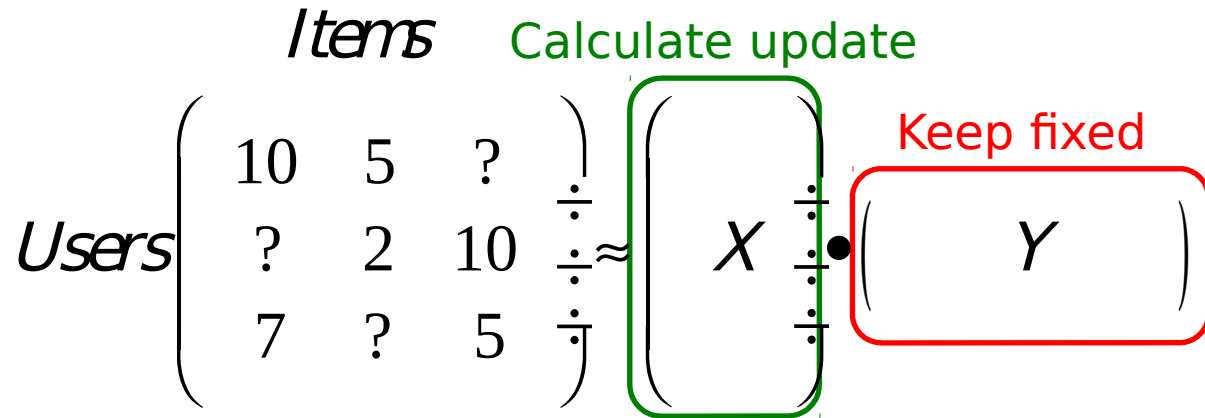
- Update item matrix

$$\begin{array}{c} \text{Users} \\ \left(\begin{array}{ccc} 10 & 5 & ? \\ ? & 2 & 10 \\ 7 & ? & 5 \end{array} \right) \end{array} \begin{array}{c} \text{Items} \\ \left(\begin{array}{c} \vdots \\ \vdots \\ \vdots \end{array} \right) \end{array} \approx \begin{array}{c} \text{Keep fixed} \\ \left(\begin{array}{c} \vdots \\ X \\ \vdots \end{array} \right) \end{array} \bullet \begin{array}{c} \text{Calculate update} \\ \left(\begin{array}{c} \vdots \\ Y \\ \vdots \end{array} \right) \end{array}$$

ALS Algorithm contd.

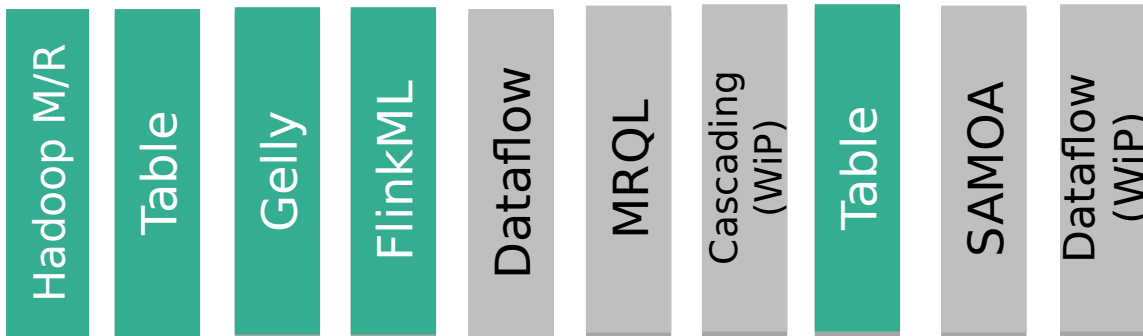


- Repeat update step until convergence



Apache Flink


What is Apache Flink?



Apache Flink deep-dive
by Stephan Ewen
Tomorrow, 12:20 - 13:00
on Stage 3

DataSet (Java/Scala/Python)

DataStream (Java/Scala)



Local

Remote

Yarn

Tez

Embedded

Why Using Flink for ALS?



- Expressive API
- Pipelined stream processor
- Closed loop iterations
- Operations on managed memory

Expressive APIs



- DataSet: Abstraction for distributed data
- Computation specified as sequence of lazily evaluated transformations

```
case class Word(word: String, frequency: Int)
```

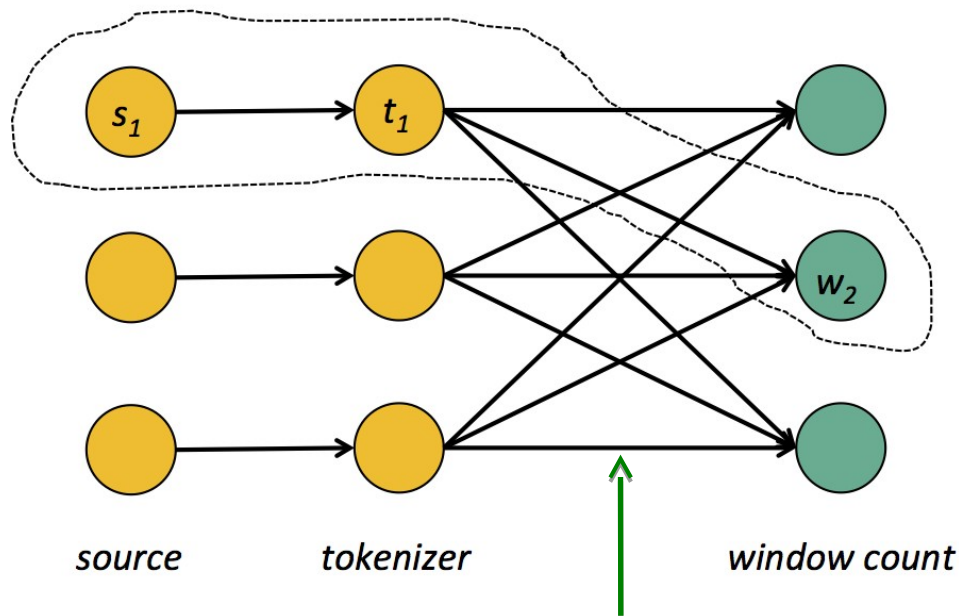
```
val lines: DataSet[String] = env.readTextFile(...)
```

```
lines.flatMap(line => line.split(" ").map(word => Word(word, 1))  
  .groupBy("word").sum("frequency")  
  .print())
```

Pipelined Stream Processor

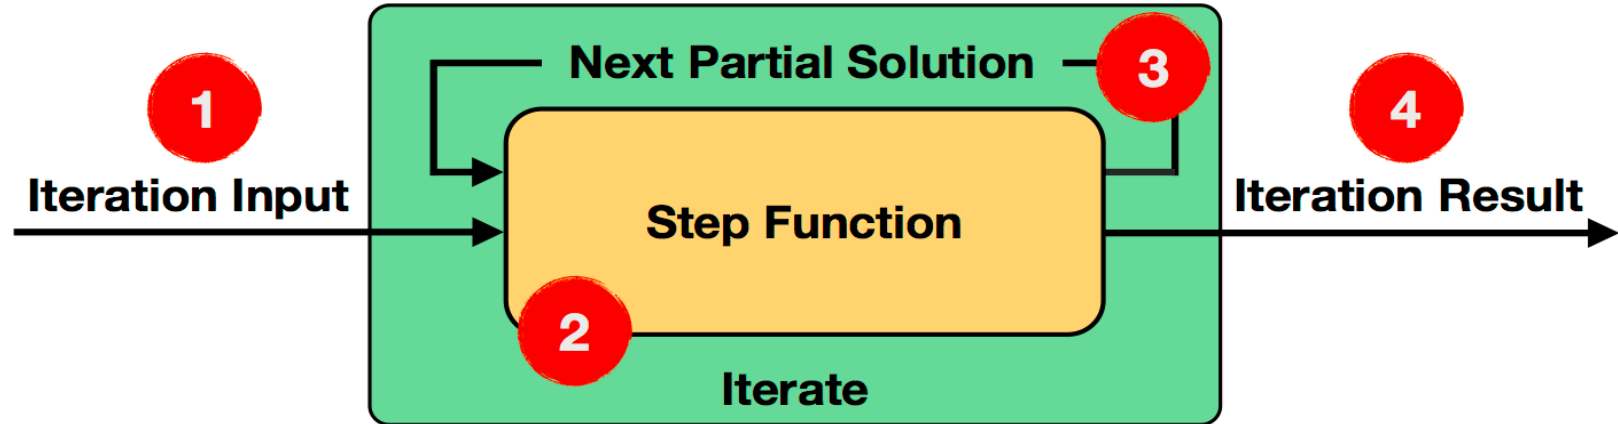


*Complete pipeline online
concurrently*



Avoiding materialization of
intermediate results


Iterate in the Dataflow



Memory Management




Custom
Data Objects



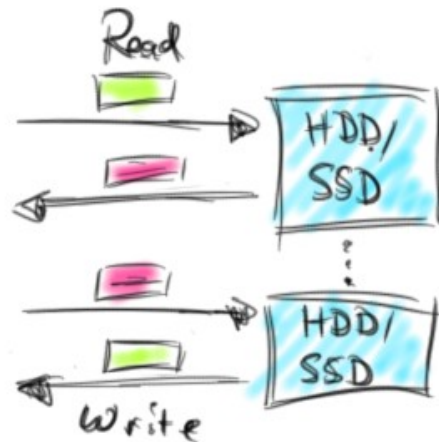
Efficient
De/Serialization



Managed Memory
32KB memory segments



Destage to Local FS
at Need

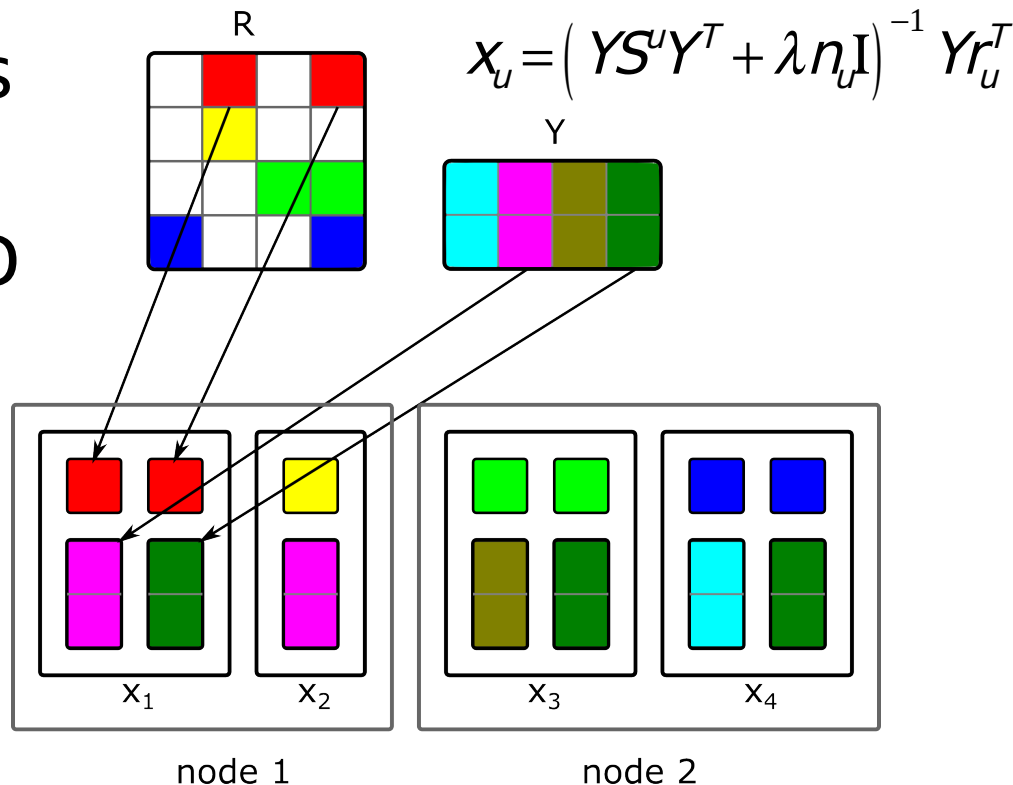


ALS implementations with Apache Flink

Naïve Implementation



1. Join item vectors with ratings
2. Group on user ID
3. Compute new user vectors



Pros and Cons of Naïve ALS

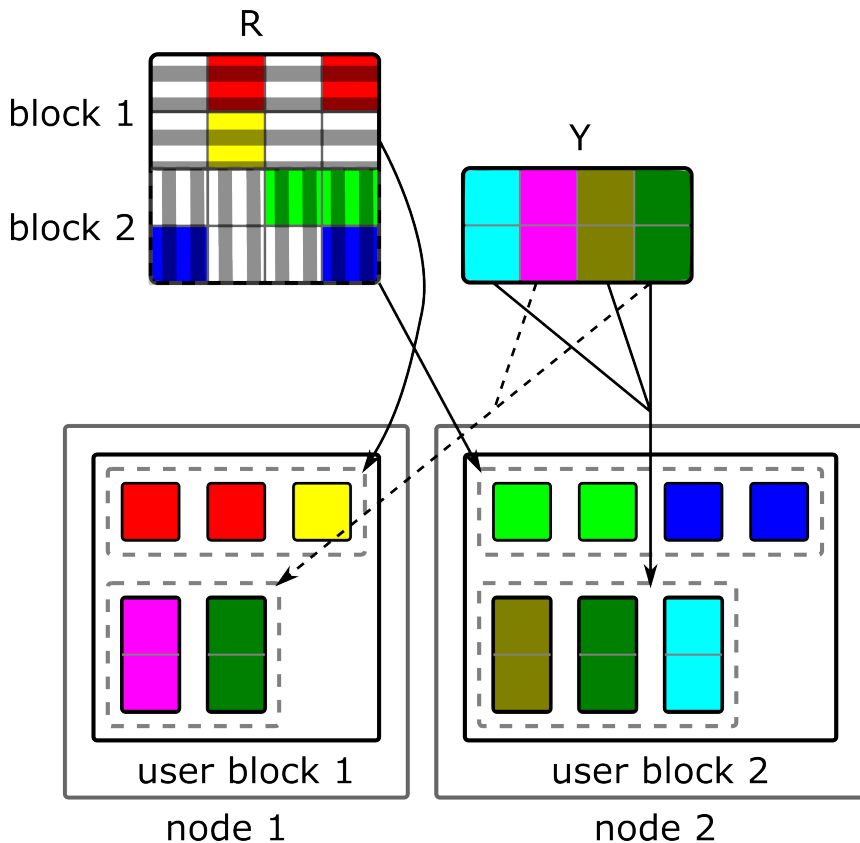


- Pros
 - Easy to implement
- Cons
 - Item vectors are sent redundantly to network nodes
 - Two shuffle steps make execution expensive

Blocked ALS Implementation



1. Create user and item rating blocks
2. Cache them on worker nodes
3. Send all item vectors needed by user rating block bundled
4. Compute block of item vectors



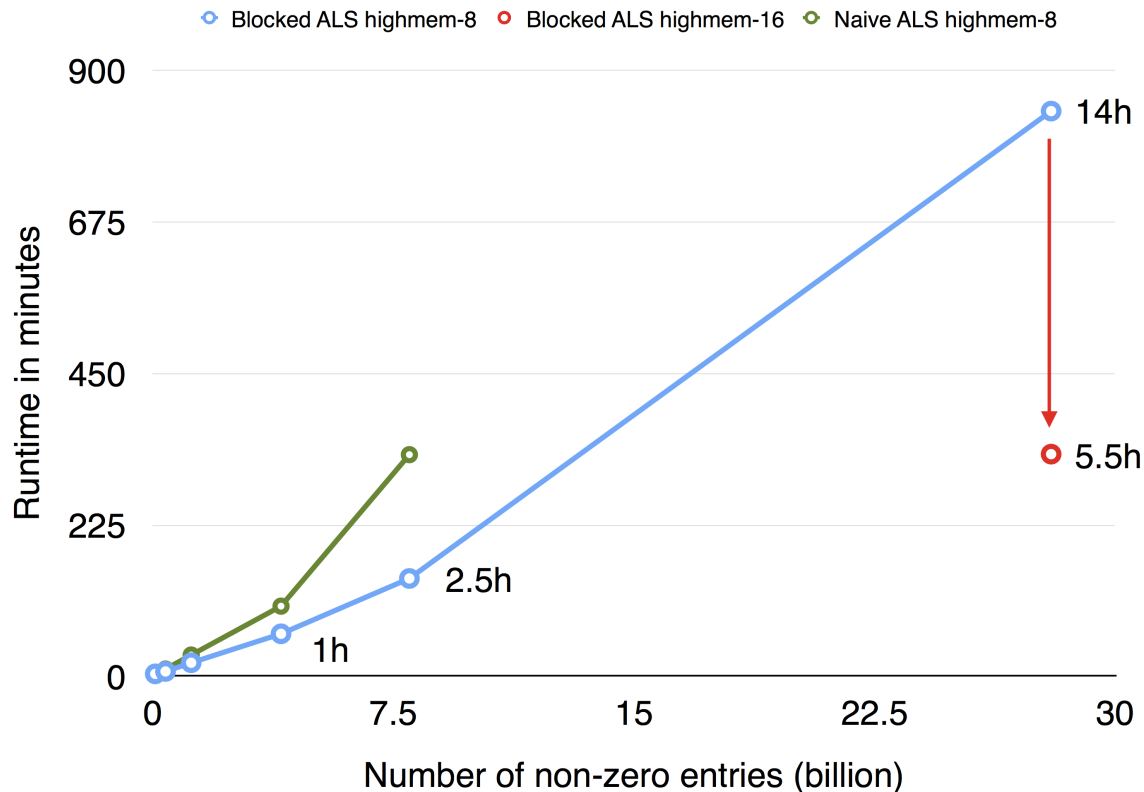
Based on Spark's MLlib implementation

Pros and Cons of Blocked ALS



- Pros
 - Reduces network load by avoiding data duplication
 - Caching ratings: Only one shuffle step needed
- Cons
 - Duplicates the rating matrix (user block/item block partitioning)

Performance Comparison



- 40 node GCE cluster, highmem-8
- 10 ALS iterations with 50 latent factors
- Rating matrix has 28 billion non zero entries: **Scale of Netflix or Spotify**

Machine Learning with FlinkML



- FlinkML contains blocked ALS
- Support for many other tasks
 - Clustering
 - Regression
 - Classification
- scikit-learn like pipeline support

```
val als = ALS()

val ratingDS =
  env.readCsvFile[(Int, Int, Double)](ratingData)

val parameters = ParameterMap()
  .add(ALS.Iterations, 10)
  .add(ALS.NumFactors, 50)
  .add(ALS.Lambda, 1.5)

als.fit(ratingDS, parameters)

val testingDS = env.readCsvFile[(Int, Int)](testingData)

val predictions = als.predict(testingDS)
```

Closing

What Have You Seen?



- How to use collaborative filtering to make recommendations
- Apache Flink, a powerful parallel stream processing engine
- How to use Apache Flink and alternating least squares to factorize really large matrices



Flink *Forward*

BERLIN 12/13 OCT 2015



flink.apache.org
@ApacheFlink