



Music Recommendations at Scale with Spark

Chris Johnson
@MrChrisJohnson



June 27, 2014

Who am I??

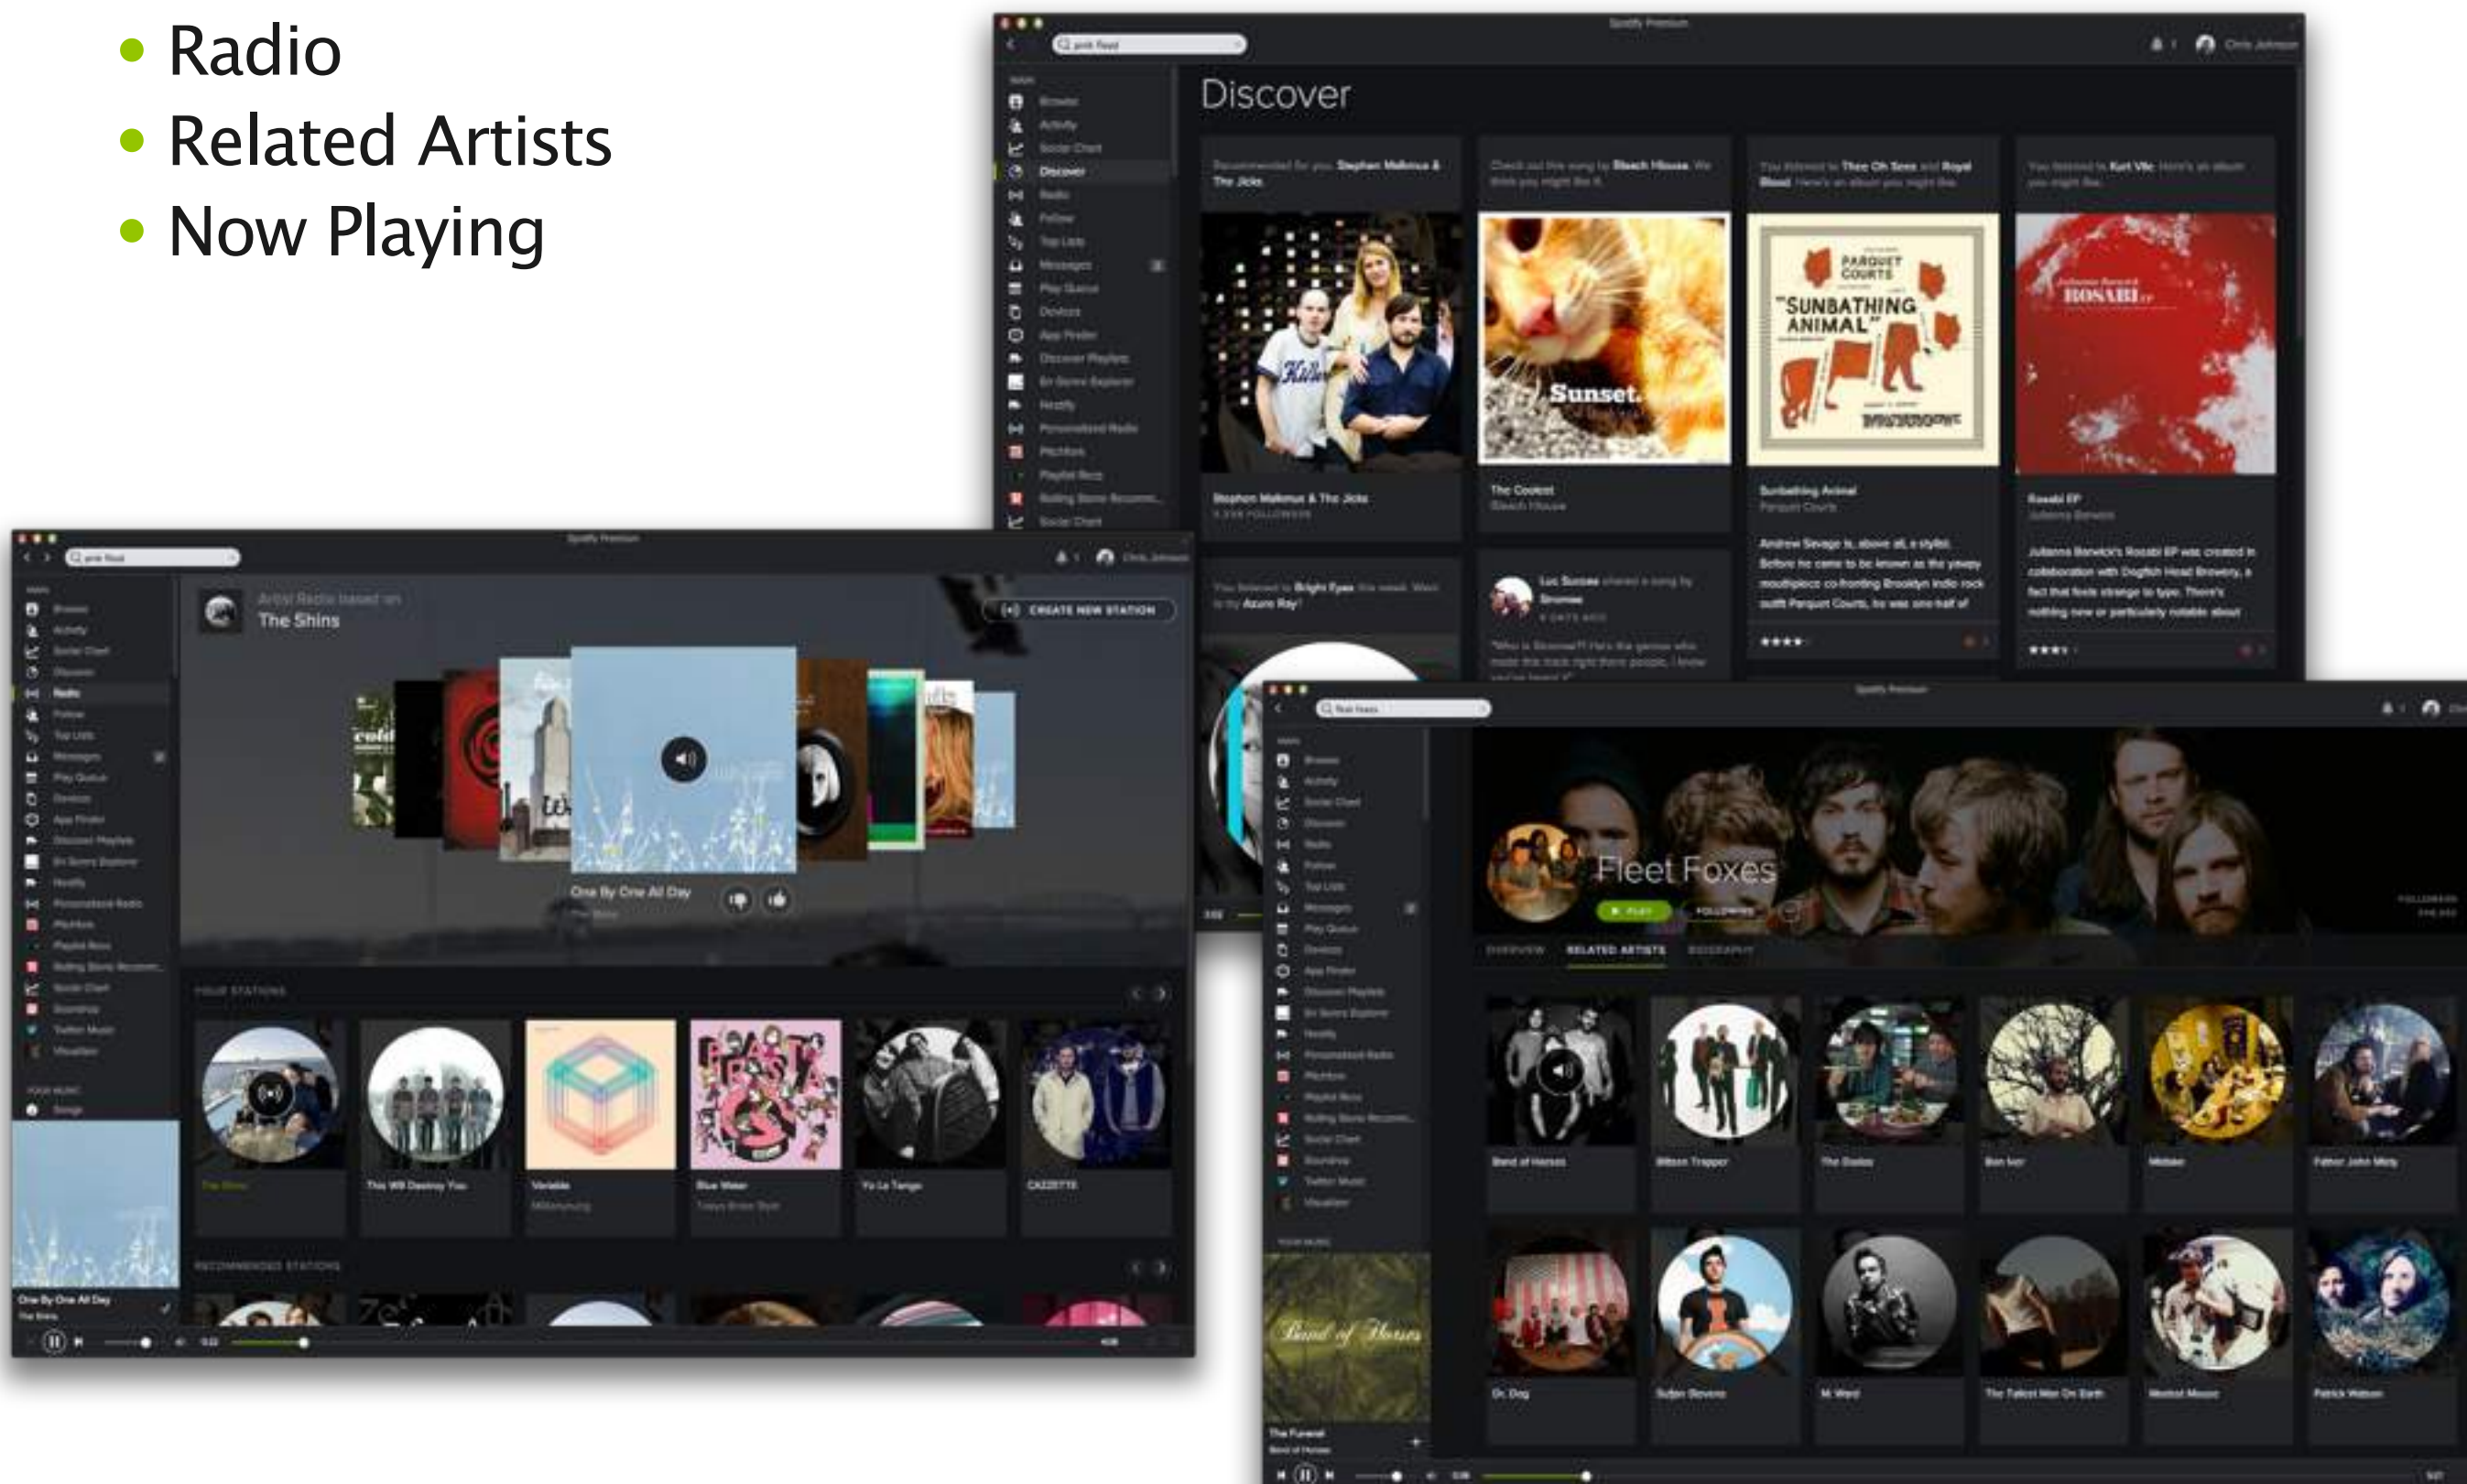
- Chris Johnson
 - Machine Learning guy from NYC
 - Focused on music recommendations
 - Formerly a PhD student at **UT Austin**



Recommendations at Spotify

3

- Discover (personalized recommendations)
- Radio
- Related Artists
- Now Playing



How can we find good recommendations?

- Manual Curation



- Manually Tag Attributes



- Audio Content, Metadata, Text Analysis



- Collaborative Filtering



How can we find good recommendations?

- Manual Curation



- Manually Tag Attributes



- Audio Content, Metadata, Text Analysis



- Collaborative Filtering



Collaborative Filtering – “The Netflix Prize”⁶



Collaborative Filtering

7





Image via Erik Bernhardsson



Explicit Matrix Factorization

- Approximate ratings matrix by the product of low-dimensional user and movie matrices
- Minimize RMSE (root mean squared error)



$$\min_{x,y} \sum_{u,i} (r_{ui} - x_u^T y_i - \beta_u - \beta_i)^2 + \lambda (\sum_u \|x_u\|^2 + \sum_i \|y_i\|^2)$$

- r_{ui} = user u 's rating for movie i
- x_u = user u 's latent factor vector
- x_i = item i 's latent factor vector
- β_u = bias for user u
- β_i = bias for item i
- λ = regularization parameter

Implicit Matrix Factorization

11

- Instead of explicit ratings use binary labels
 - 1 = streamed, 0 = never streamed
- Minimize weighted RMSE (root mean squared error) using a function of total streams as weights

$$\text{Users} \begin{pmatrix} 1 & 0 & 0 & 0 & 1 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 & 0 & 1 & 0 & 0 \\ 1 & 0 & 1 & 0 & 0 & 0 & 1 & 1 \\ 0 & 1 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 & 1 & 0 & 0 & 1 \end{pmatrix} \approx \underbrace{\begin{pmatrix} X \\ Y \end{pmatrix}}_f$$

$$\min_{x,y} \sum_{u,i} c_{ui} (p_{ui} - x_u^T y_i - \beta_u - \beta_i)^2 + \lambda (\sum_u \|x_u\|^2 + \sum_i \|y_i\|^2)$$

- p_{ui} = 1 if user u streamed track i else 0
- $c_{ui} = 1 + \alpha r_{ui}$
- x_u = user u 's latent factor vector
- x_i = item i 's latent factor vector
- β_u = bias for user u
- β_i = bias for item i
- λ = regularization parameter

Alternating Least Squares (ALS)

12

- Instead of explicit ratings use binary labels
 - 1 = streamed, 0 = never streamed
- Minimize weighted RMSE (root mean squared error) using a function of total streams as weights

$$\begin{array}{c} \text{Users} \end{array} \begin{pmatrix} 1 & 0 & 0 & 0 & 1 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 & 0 & 1 & 0 & 0 \\ 1 & 0 & 1 & 0 & 0 & 0 & 1 & 1 \\ 0 & 1 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 & 1 & 0 & 0 & 1 \end{pmatrix} \approx \underbrace{\begin{pmatrix} X \end{pmatrix}}_f \left(\begin{array}{c} Y \end{array} \right) f$$

Songs

Fix songs

$$\min_{x,y} \sum_{u,i} c_{ui} (p_{ui} - x_u^T y_i - \beta_u - \beta_i)^2 + \lambda (\sum_u \|x_u\|^2 + \sum_i \|y_i\|^2)$$

- p_{ui} = 1 if user u streamed track i else 0
- $c_{ui} = 1 + \alpha r_{ui}$
- x_u = user u 's latent factor vector
- x_i = item i 's latent factor vector
- β_u = bias for user u
- β_i = bias for item i
- λ = regularization parameter

Alternating Least Squares (ALS)

13

- Instead of explicit ratings use binary labels
 - 1 = streamed, 0 = never streamed
- Minimize weighted RMSE (root mean squared error) using a function of total streams as weights

$$\text{Users} \begin{pmatrix} 1 & 0 & 0 & 0 & 1 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 & 0 & 1 & 0 & 0 \\ 1 & 0 & 1 & 0 & 0 & 0 & 1 & 1 \\ 0 & 1 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 & 1 & 0 & 0 & 1 \end{pmatrix} \approx \underbrace{\begin{pmatrix} \vdots \\ \vdots \\ \vdots \end{pmatrix}}_f X \left(\begin{pmatrix} \vdots \\ \vdots \\ \vdots \end{pmatrix} Y \right) f$$

$x_u = (Y^T C^u Y + \lambda I)^{-1} Y^T C^u p(u)$

Solve for users

Fix songs

$$\min_{x,y} \sum_{u,i} c_{ui} (p_{ui} - x_u^T y_i - \beta_u - \beta_i)^2 + \lambda (\sum_u \|x_u\|^2 + \sum_i \|y_i\|^2)$$

- p_{ui} = 1 if user u streamed track i else 0
- $c_{ui} = 1 + \alpha r_{ui}$
- x_u = user u 's latent factor vector
- x_i = item i 's latent factor vector
- β_u = bias for user u
- β_i = bias for item i
- λ = regularization parameter

Alternating Least Squares (ALS)

14

- Instead of explicit ratings use binary labels
 - 1 = streamed, 0 = never streamed
- Minimize weighted RMSE (root mean squared error) using a function of total streams as weights

$$\text{Users} \begin{pmatrix} 1 & 0 & 0 & 0 & 1 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 & 0 & 1 & 0 & 0 \\ 1 & 0 & 1 & 0 & 0 & 0 & 1 & 1 \\ 0 & 1 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 & 1 & 0 & 0 & 1 \end{pmatrix} \approx \underbrace{\begin{pmatrix} X \\ \vdots \\ X \end{pmatrix}}_f (Y) \} f$$

Songs

Fix users

$$\min_{x,y} \sum_{u,i} c_{ui} (p_{ui} - x_u^T y_i - \beta_u - \beta_i)^2 + \lambda (\sum_u \|x_u\|^2 + \sum_i \|y_i\|^2)$$

- p_{ui} = 1 if user u streamed track i else 0
- $c_{ui} = 1 + \alpha r_{ui}$
- x_u = user u 's latent factor vector
- x_i = item i 's latent factor vector
- β_u = bias for user u
- β_i = bias for item i
- λ = regularization parameter

Alternating Least Squares (ALS)

15

- Instead of explicit ratings use binary labels
 - 1 = streamed, 0 = never streamed
- Minimize weighted RMSE (root mean squared error) using a function of total streams as weights

Users

$$\begin{pmatrix} 1 & 0 & 0 & 0 & 1 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 & 0 & 1 & 0 & 0 \\ 1 & 0 & 1 & 0 & 0 & 0 & 1 & 1 \\ 0 & 1 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 & 1 & 0 & 0 & 1 \end{pmatrix} \approx \underbrace{\begin{pmatrix} \vdots \\ \vdots \\ \vdots \\ \vdots \\ \vdots \\ \vdots \end{pmatrix}}_f X \left(\begin{pmatrix} \vdots \\ \vdots \\ \vdots \\ \vdots \\ \vdots \\ \vdots \end{pmatrix} Y \right) f$$

Songs

Solve for songs

Fix users $y_i = (X^T C^i X + \lambda I)^{-1} X^T C^i p(i)$

$$\min_{x,y} \sum_{u,i} c_{ui} (p_{ui} - x_u^T y_i - \beta_u - \beta_i)^2 + \lambda (\sum_u \|x_u\|^2 + \sum_i \|y_i\|^2)$$

- p_{ui} = 1 if user u streamed track i else 0
- $c_{ui} = 1 + \alpha r_{ui}$
- x_u = user u 's latent factor vector
- x_i = item i 's latent factor vector
- β_u = bias for user u
- β_i = bias for item i
- λ = regularization parameter

Alternating Least Squares (ALS)

16

- Instead of explicit ratings use binary labels
 - 1 = streamed, 0 = never streamed
- Minimize weighted RMSE (root mean squared error) using a function of total streams as weights

Users $\begin{pmatrix} 1 & 0 & 0 & 0 & 1 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 & 0 & 1 & 0 & 0 \\ 1 & 0 & 1 & 0 & 0 & 0 & 1 & 1 \\ 0 & 1 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 & 1 & 0 & 0 & 1 \end{pmatrix} \approx \underbrace{\begin{pmatrix} \\ \\ \\ \\ \\ \end{pmatrix}}_f X \underbrace{\left(\begin{pmatrix} \\ \\ \\ \\ \\ \end{pmatrix} Y \right)}_f \} f$

Songs

Repeat until convergence...

Solve for songs

Fix users $y_i = (X^T C^i X + \lambda I)^{-1} X^T C^i p(i)$

$$\min_{x,y} \sum_{u,i} c_{ui} (p_{ui} - x_u^T y_i - \beta_u - \beta_i)^2 + \lambda (\sum_u \|x_u\|^2 + \sum_i \|y_i\|^2)$$

- p_{ui} = 1 if user u streamed track i else 0
- $c_{ui} = 1 + \alpha r_{ui}$
- x_u = user u 's latent factor vector
- x_i = item i 's latent factor vector
- β_u = bias for user u
- β_i = bias for item i
- λ = regularization parameter

Alternating Least Squares (ALS)

17

- Instead of explicit ratings use binary labels
 - 1 = streamed, 0 = never streamed
- Minimize weighted RMSE (root mean squared error) using a function of total streams as weights

Users $\begin{pmatrix} 1 & 0 & 0 & 0 & 1 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 & 0 & 1 & 0 & 0 \\ 1 & 0 & 1 & 0 & 0 & 0 & 1 & 1 \\ 0 & 1 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 & 1 & 0 & 0 & 1 \end{pmatrix} \approx \underbrace{\begin{pmatrix} \vdots \\ \vdots \\ \vdots \\ \vdots \\ \vdots \\ \vdots \end{pmatrix}}_f X \underbrace{\left(\begin{pmatrix} \vdots \\ \vdots \\ \vdots \\ \vdots \\ \vdots \\ \vdots \end{pmatrix} \right)}_f Y \}$

Songs

Repeat until convergence...

Solve for songs

Fix users $y_i = (X^T C^i X + \lambda I)^{-1} X^T C^i p(i)$

$y_i = (X^T X + X^T (C^i - I) X + \lambda I)^{-1} X^T C^i p(i)$

- $p_{ui} = 1$ if user u streamed track i else 0
- $c_{ui} = 1 + \alpha r_{ui}$
- x_u = user u 's latent factor vector
- x_i = item i 's latent factor vector
- β_u = bias for user u
- β_i = bias for item i
- λ = regularization parameter

Alternating Least Squares

18


```
2 def als_iteration(user, counts, solve_vecs, fixed_vecs, num_factors=40, reg_param=0.8):
3     '''
4         @param user:      True if solving for user vectors
5         @param counts:    scipy.sparse matrix containing implicit
6                           user-item counts * alpha
7         @param solve_vecs: scipy.sparse vector of latent factors you
8                           wish to solve for
9         @param fixed_vecs: scipy.sparse vector of fixed latent factors
10        @param reg_param:  regularization parameter (lambda)
11
12    '''
13    num_fixed = fixed_vecs.shape[0]
14    YTY = fixed_vecs.T.dot(fixed_vecs)
15    eye = scipy.sparse.eye(num_fixed)
16    lambda_eye = reg_param * scipy.sparse.eye(num_factors)
17
18    for i in xrange(solve_vecs.shape[0]):
19        if user:
20            counts_i = counts[i].toarray()
21        else:
22            counts_i = counts[:, i].T.toarray()
23        CuI = scipy.sparse.diags(counts_i, [0])
24        pu = counts_i.copy()
25        pu[numpy.where(pu != 0)] = 1.0
26        YTCuIY = fixed_vecs.T.dot(CuI).dot(fixed_vecs)
27        YTCupu = fixed_vecs.T.dot(CuI + eye).dot(scipy.sparse.csr_matrix(pu).T)
28        xu = scipy.sparse.linalg.spsolve(YTY + YTCuIY + lambda_eye, YTCupu)
29        solve_vecs[i] = xu
30
31    return solve_vecs
```

code: <https://github.com/MrChrisJohnson/implicitMF>



Scaling up Implicit Matrix Factorization with Hadoop

20



Hadoop at Spotify 2009

21



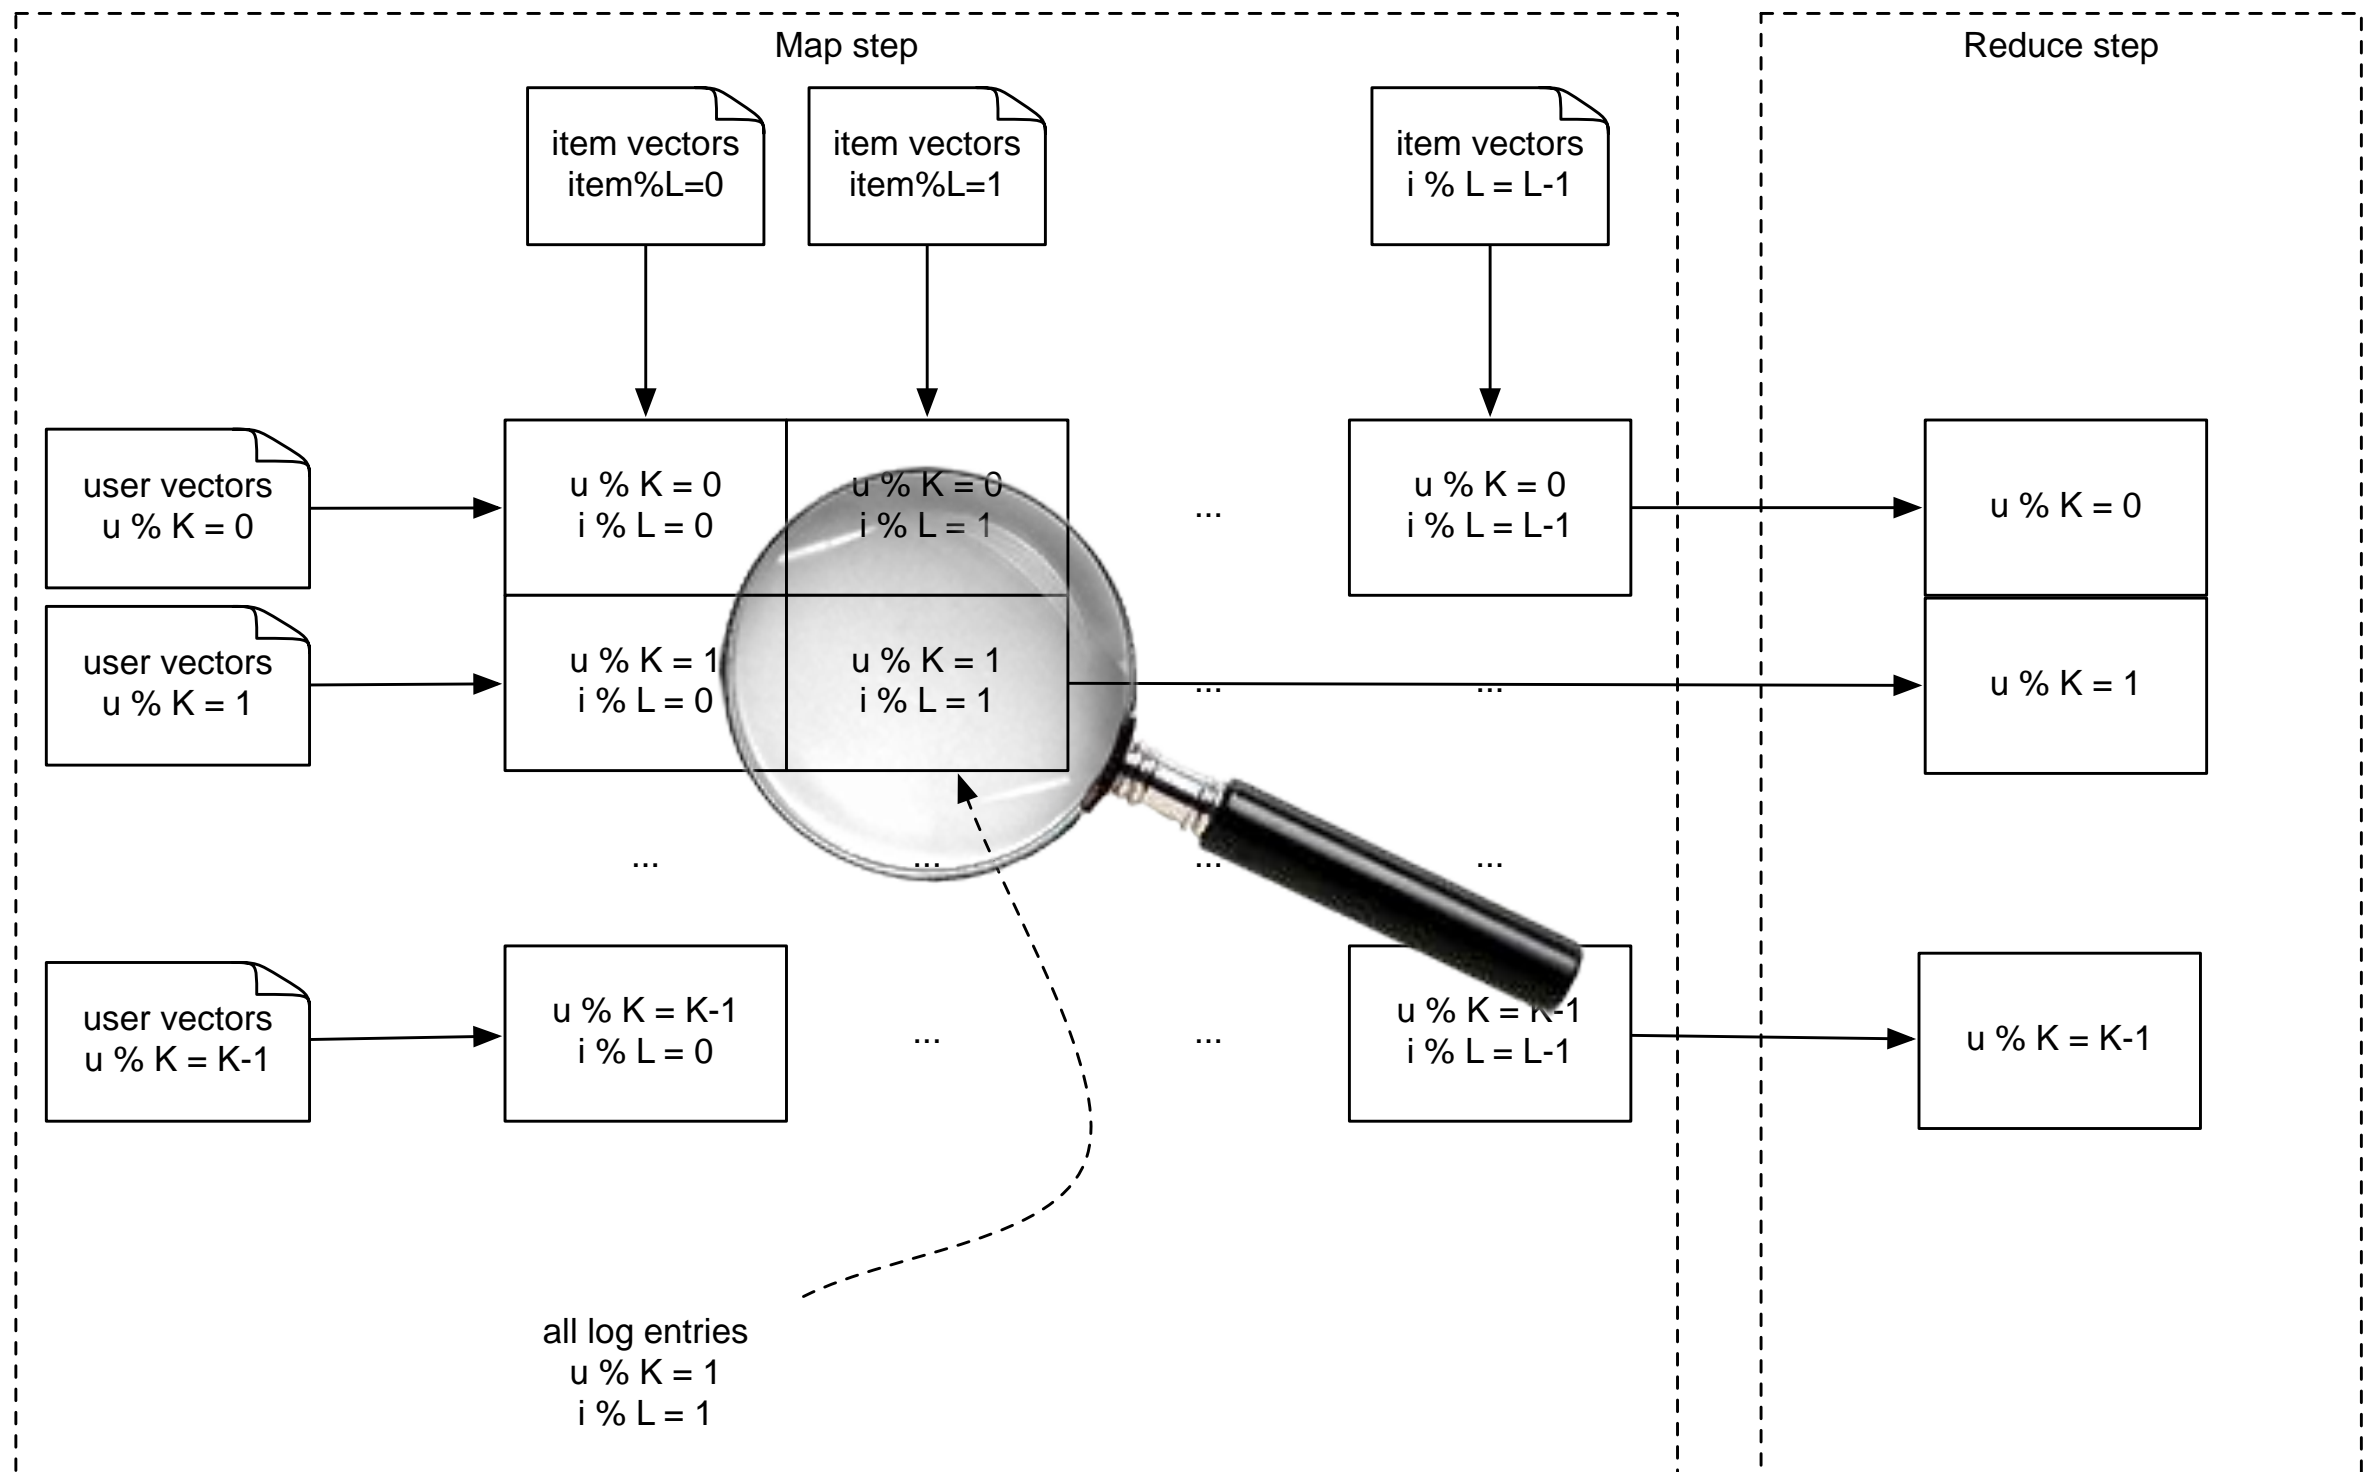
Hadoop at Spotify 2014

22

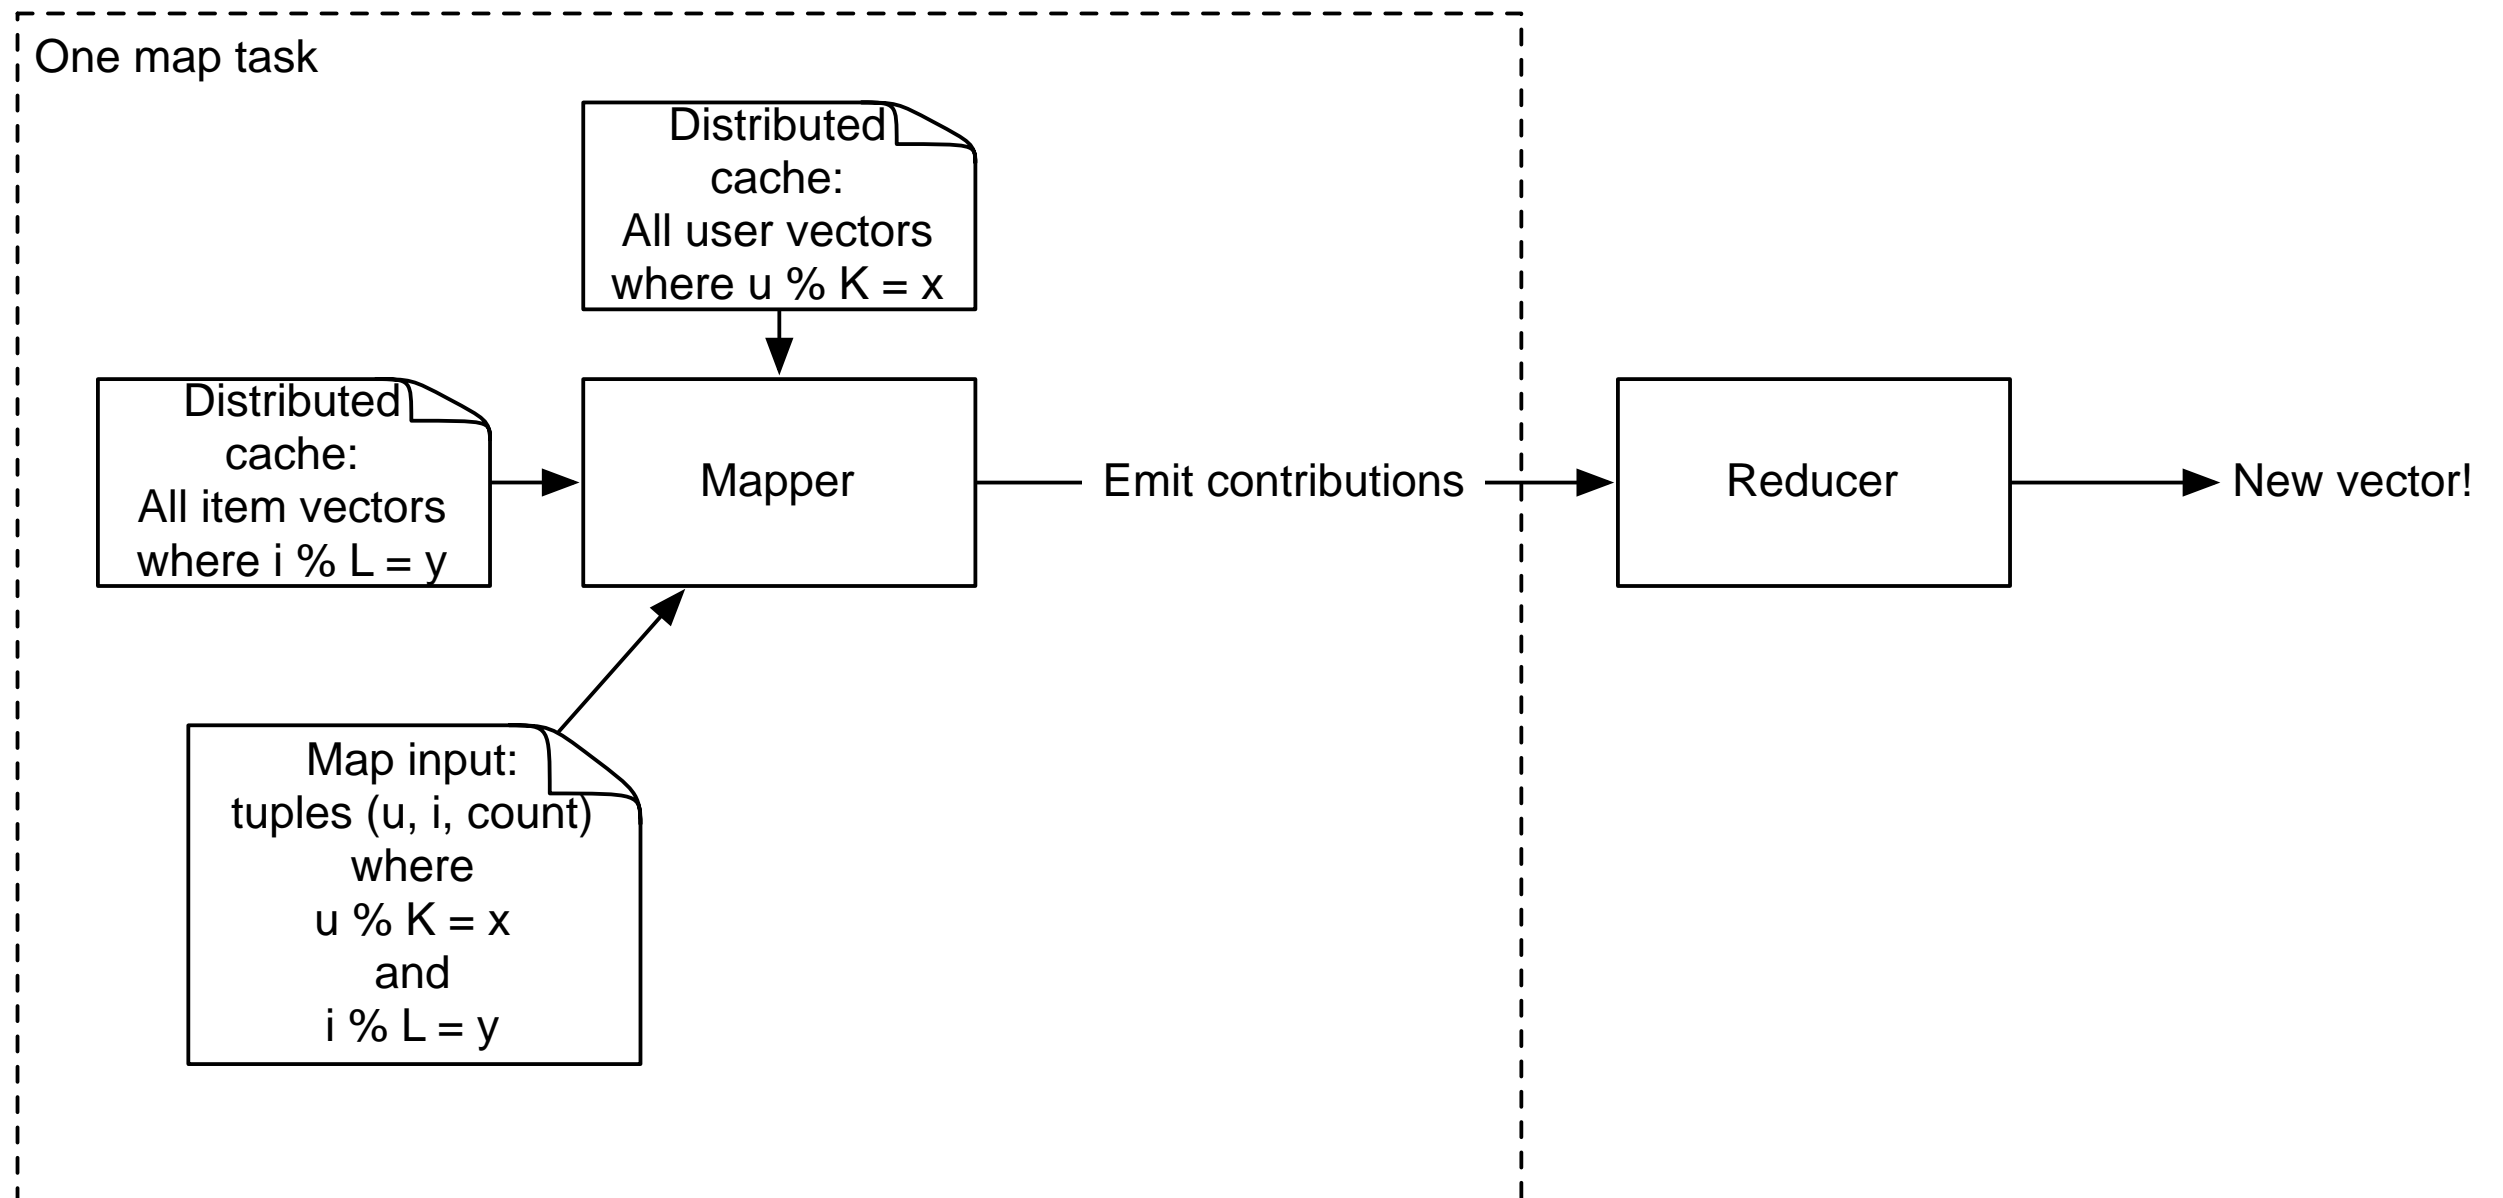
700 Nodes in our London data center




Implicit Matrix Factorization with Hadoop²³



Implicit Matrix Factorization with Hadoop²⁴




Hadoop suffers from I/O overhead




Spark to the rescue!!




26






Vs





First Attempt (broadcast everything)

- For each iteration:
 1. Compute $Y^t Y$ over item vectors and broadcast
 2. Broadcast item vectors
 3. Group ratings by user
 4. Solve for optimal user vector



worker 1




worker 2




worker 3



worker 4



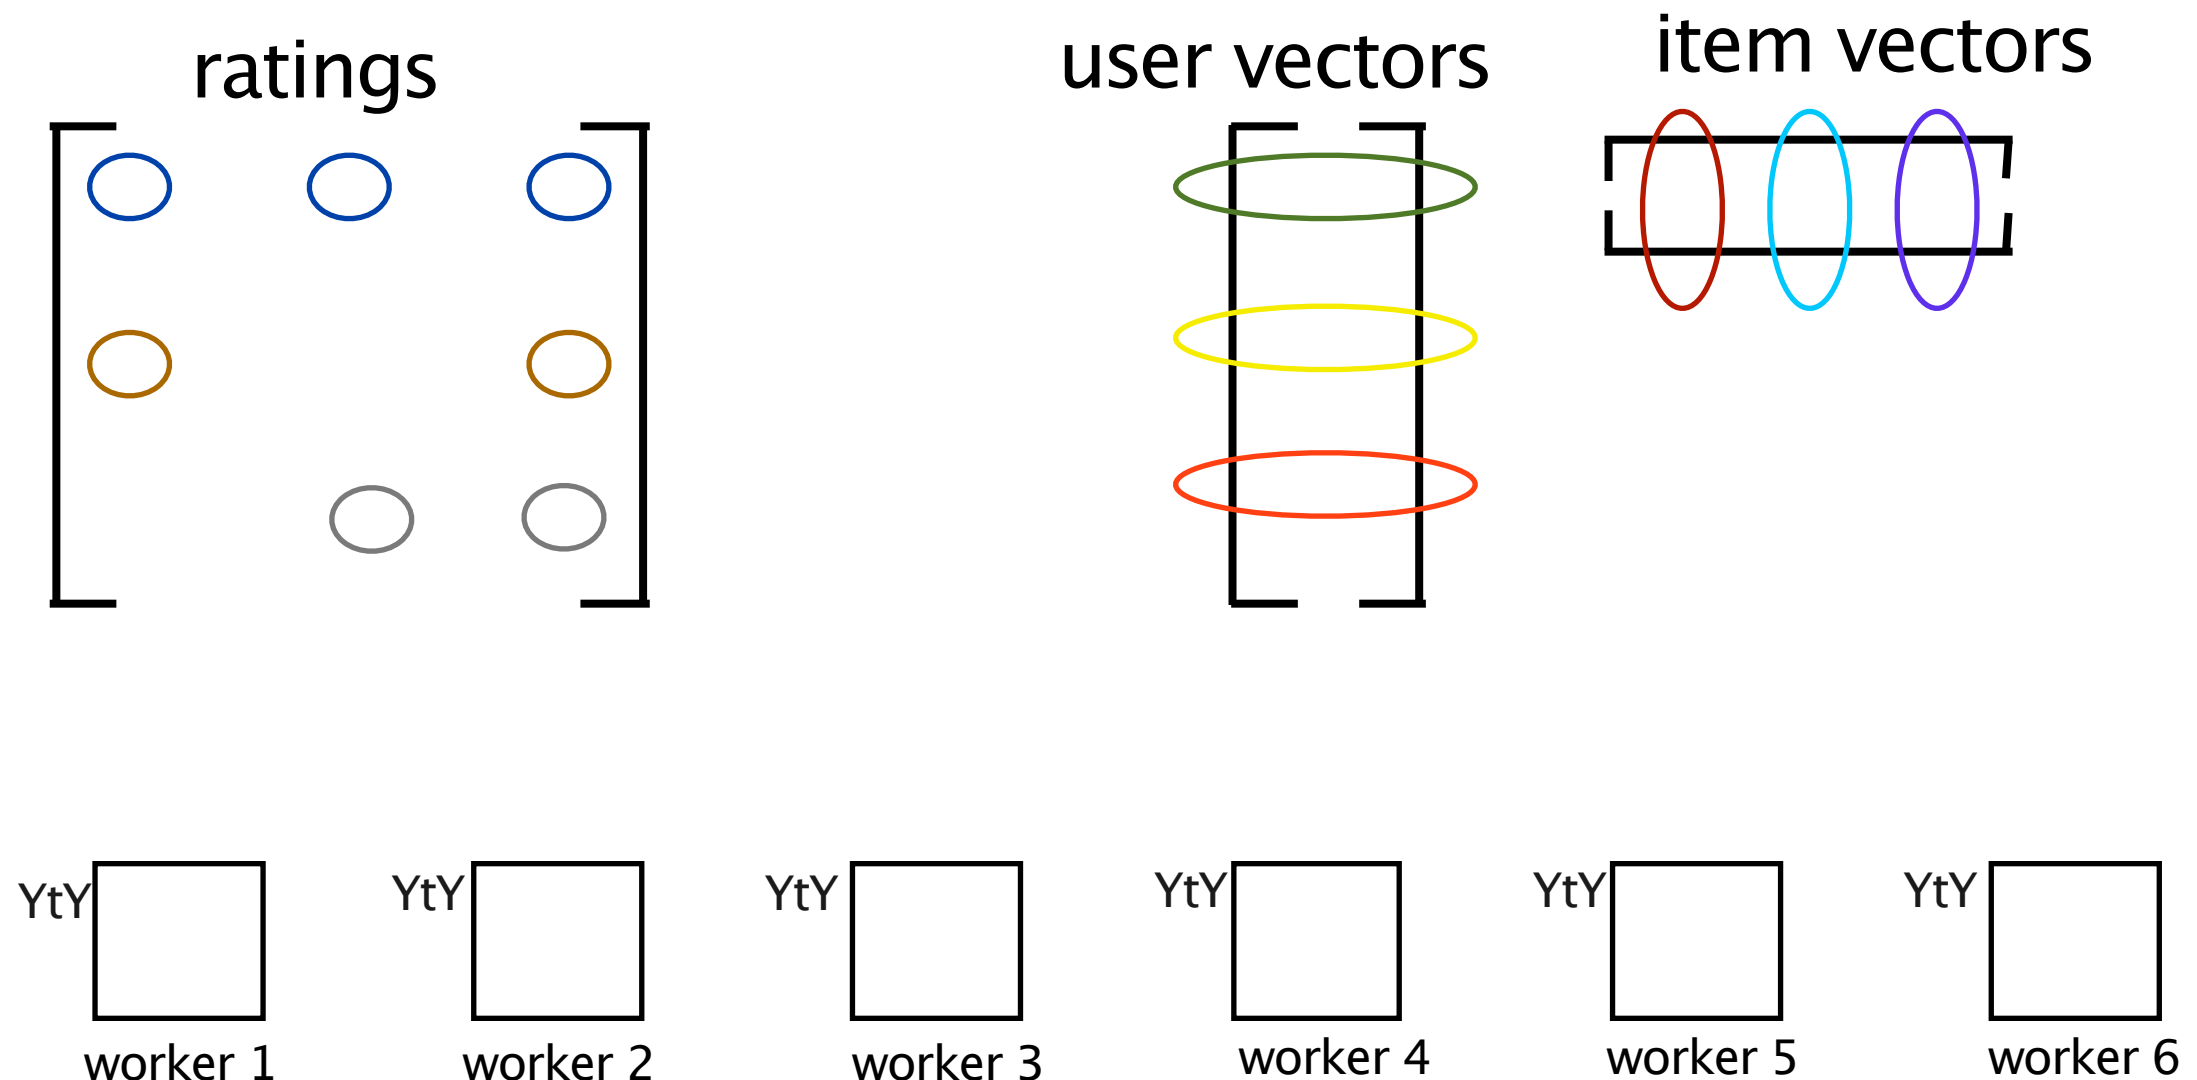
worker 5



worker 6

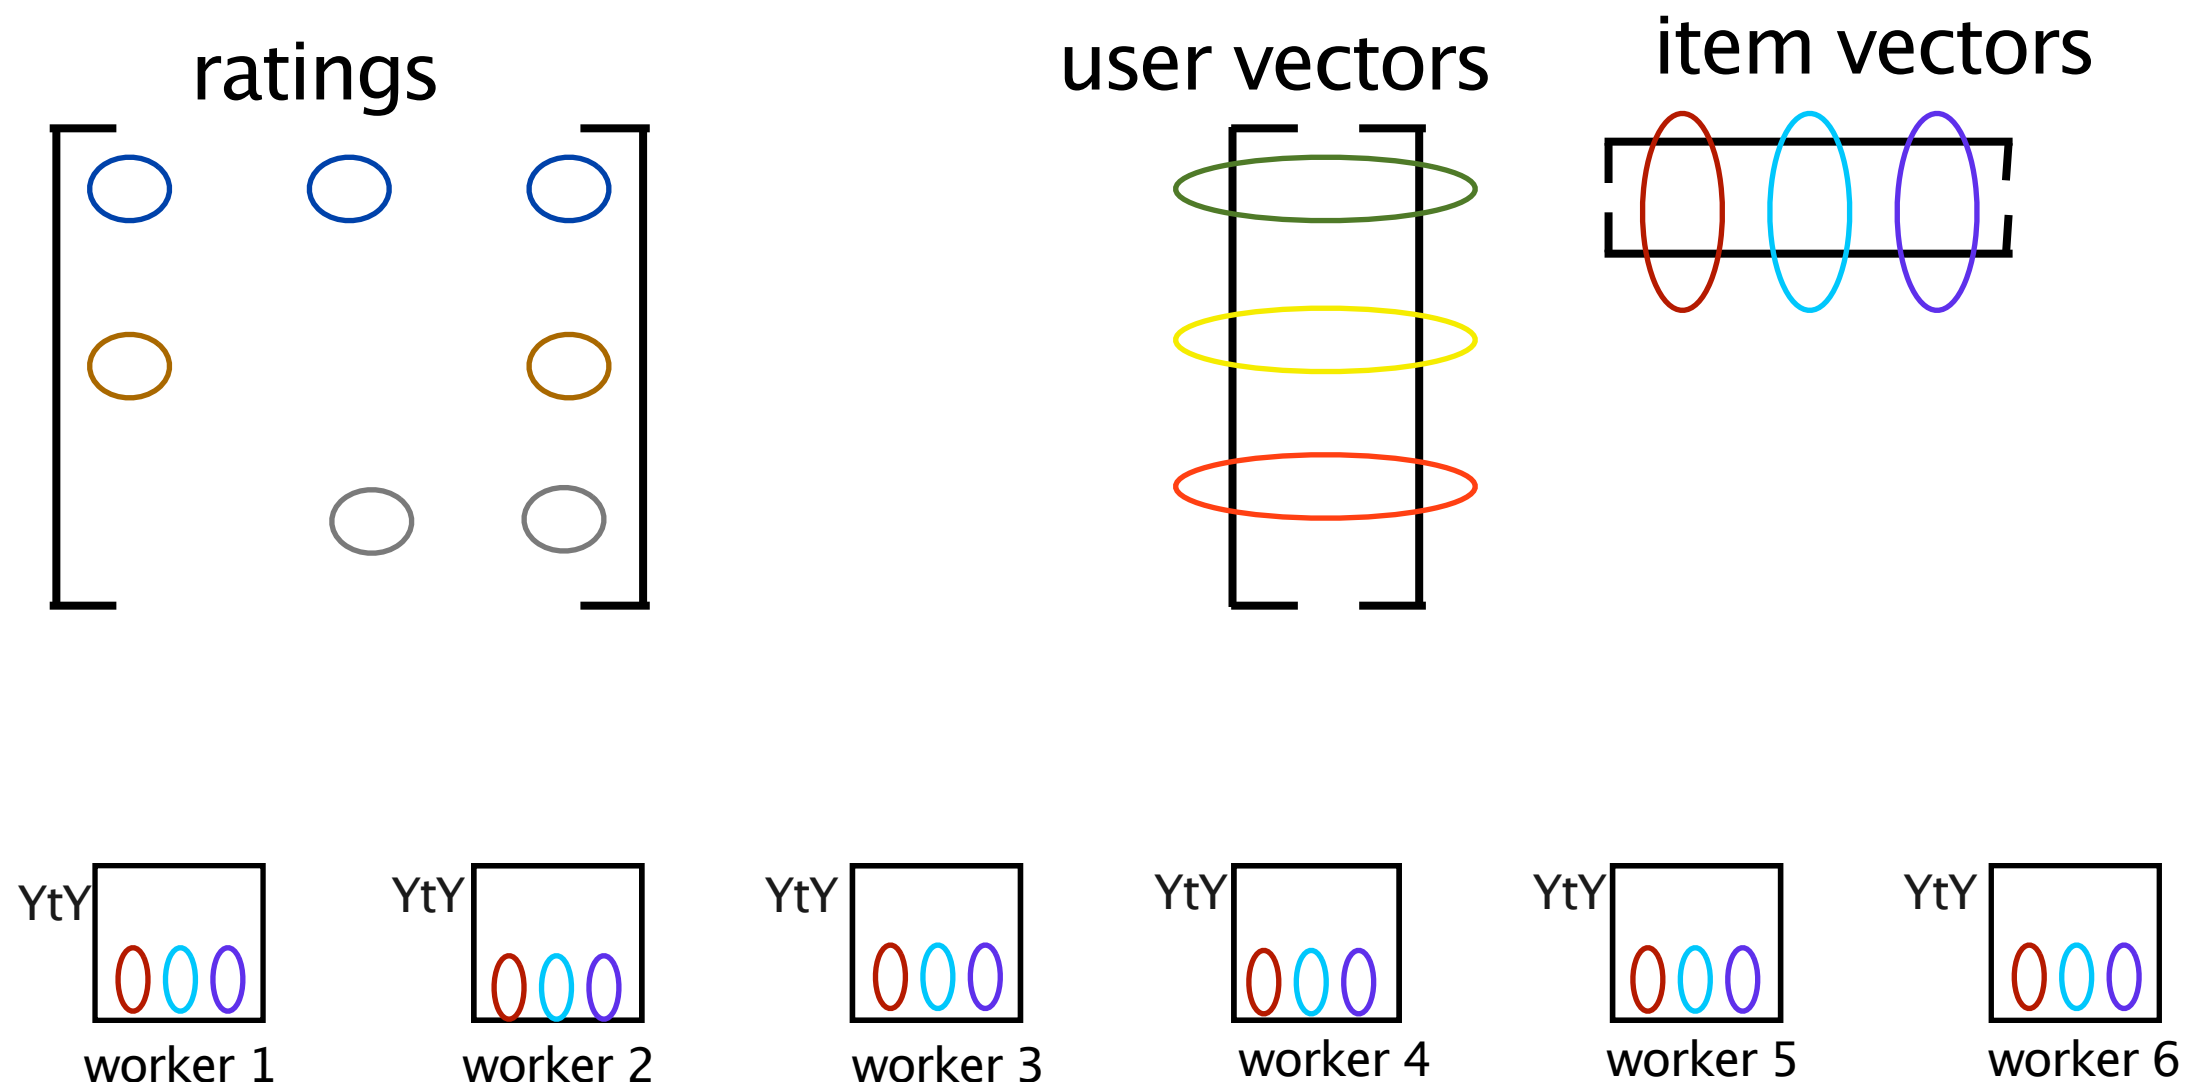
First Attempt (broadcast everything)

- For each iteration:
 1. Compute $Y^t Y$ over item vectors and broadcast
 2. Broadcast item vectors
 3. Group ratings by user
 4. Solve for optimal user vector



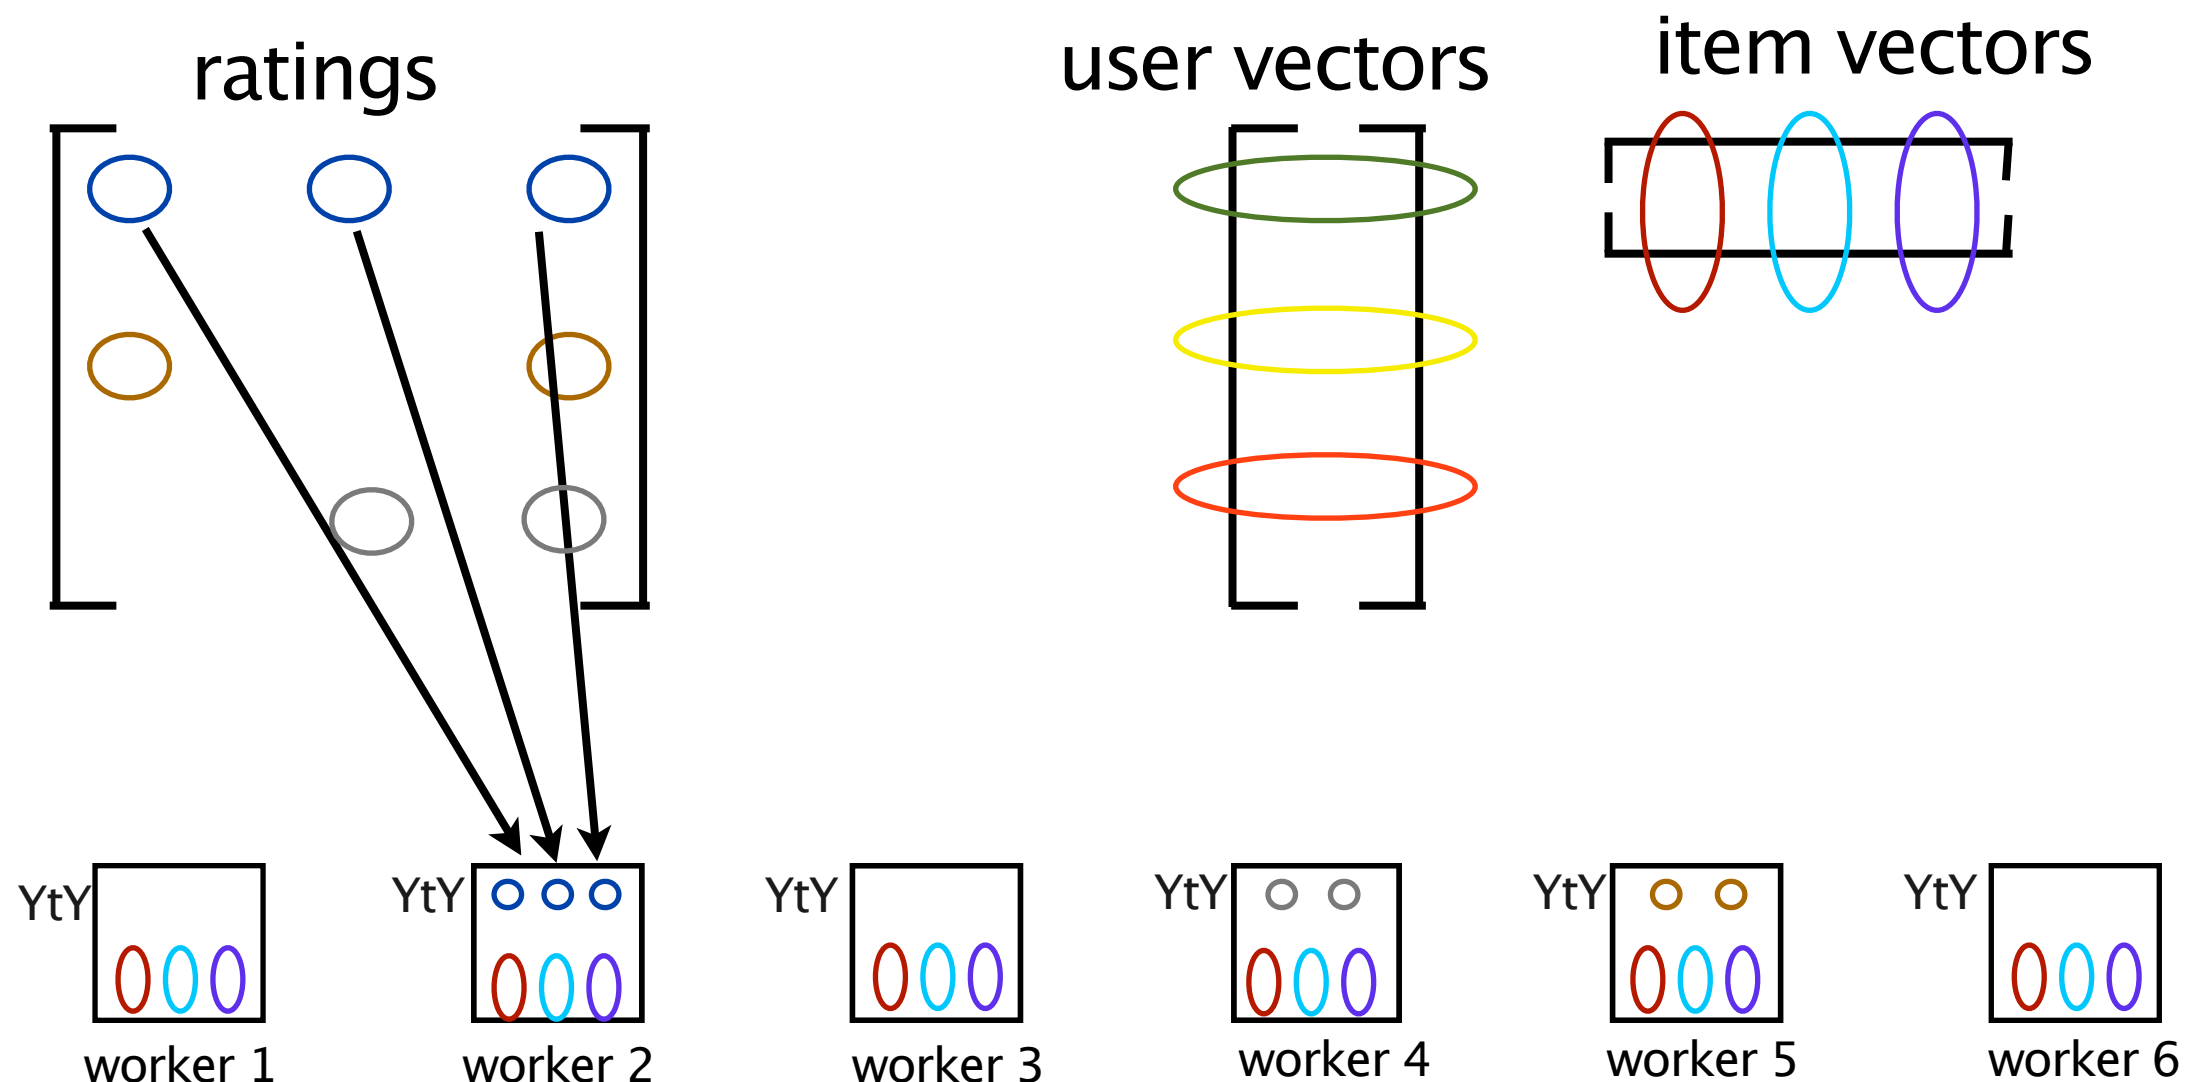
First Attempt (broadcast everything)

- For each iteration:
 1. Compute $Y^t Y$ over item vectors and broadcast
 2. Broadcast item vectors
 3. Group ratings by user
 4. Solve for optimal user vector




First Attempt (broadcast everything)

- For each iteration:
 1. Compute $Y^t Y$ over item vectors and broadcast
 2. Broadcast item vectors
 3. Group ratings by user
 4. Solve for optimal user vector




First Attempt (broadcast everything)


- For each iteration:
 1. Compute $Y^t Y$ over item vectors and broadcast
 2. Broadcast item vectors
 3. Group ratings by user
 4. Solve for optimal user vector




user vectors




item vectors




worker 1




worker 2



worker 3



worker 4



worker 5



worker 6

First Attempt (broadcast everything)


```
def ALSIteration(ratings: RDD[(Int, Int, Double)],
                 users: RDD[(Int, DenseVector[Double])],
                 items: RDD[(Int, DenseVector[Double])]) = {

  val YtY: Broadcast[DenseMatrix[Double]] = sc.broadcast(
    items
    .map{case (item: Int, vector: DenseVector[Double]) =>
      vector * vector.t
    }.reduce{(m1: DenseMatrix[Double], m2: DenseMatrix[Double]) =>
      m1 + m2}
  )

  val itemMap = sc.broadcast(
    items
    .toLocalIterator
    .toMap
  )

  ratings
    .map{case (user: Int, item: Int, rating: Double) =>
      (user, (item, rating))}
    .groupByKey
    .map{case (user: Int, ratings: Iterable[(Int, Double)]) =>
      solveVectors(user, ratings, itemMap, YtY)
    }
}
```

First Attempt (broadcast everything)



```
def ALSIteration(ratings: RDD[(Int, Int, Double)],
                 users: RDD[(Int, DenseVector[Double])],
                 items: RDD[(Int, DenseVector[Double])]) = {

  val YtY: Broadcast[DenseMatrix[Double]] = sc.broadcast(
    items
    .map{case (item: Int, vector: DenseVector[Double]) =>
      vector * vector.t}
    .reduce{(m1: DenseMatrix[Double], m2: DenseMatrix[Double]) =>
      m1 + m2}
  )

  val itemMap = sc.broadcast(
    items
    .toLocalIterator
    .toMap
  )


  ratings
    .map{case (user: Int, item: Int, rating: Double) =>
      (user, (item, rating))}
    .groupByKey
    .map{case (user: Int, ratings: Iterable[(Int, Double)]) =>
      solveVectors(user, ratings, itemMap, YtY)}
}
```

● Cons:

- Unnecessarily shuffling all data across wire each iteration.
- Not caching ratings data
- Unnecessarily sending a full copy of user/item vectors to all workers.

Second Attempt (full gridify)


- Group ratings matrix into $K \times L$, partition, and cache
- For each iteration:
 1. Compute YtY over item vectors and broadcast
 2. For each item vector send a copy to each rating block in the $\text{item} \% L$ column
 3. Compute intermediate terms for each block (partition)
 4. Group by user, aggregate intermediate terms, and solve for optimal user vector



worker 1




worker 2




worker 3



worker 4



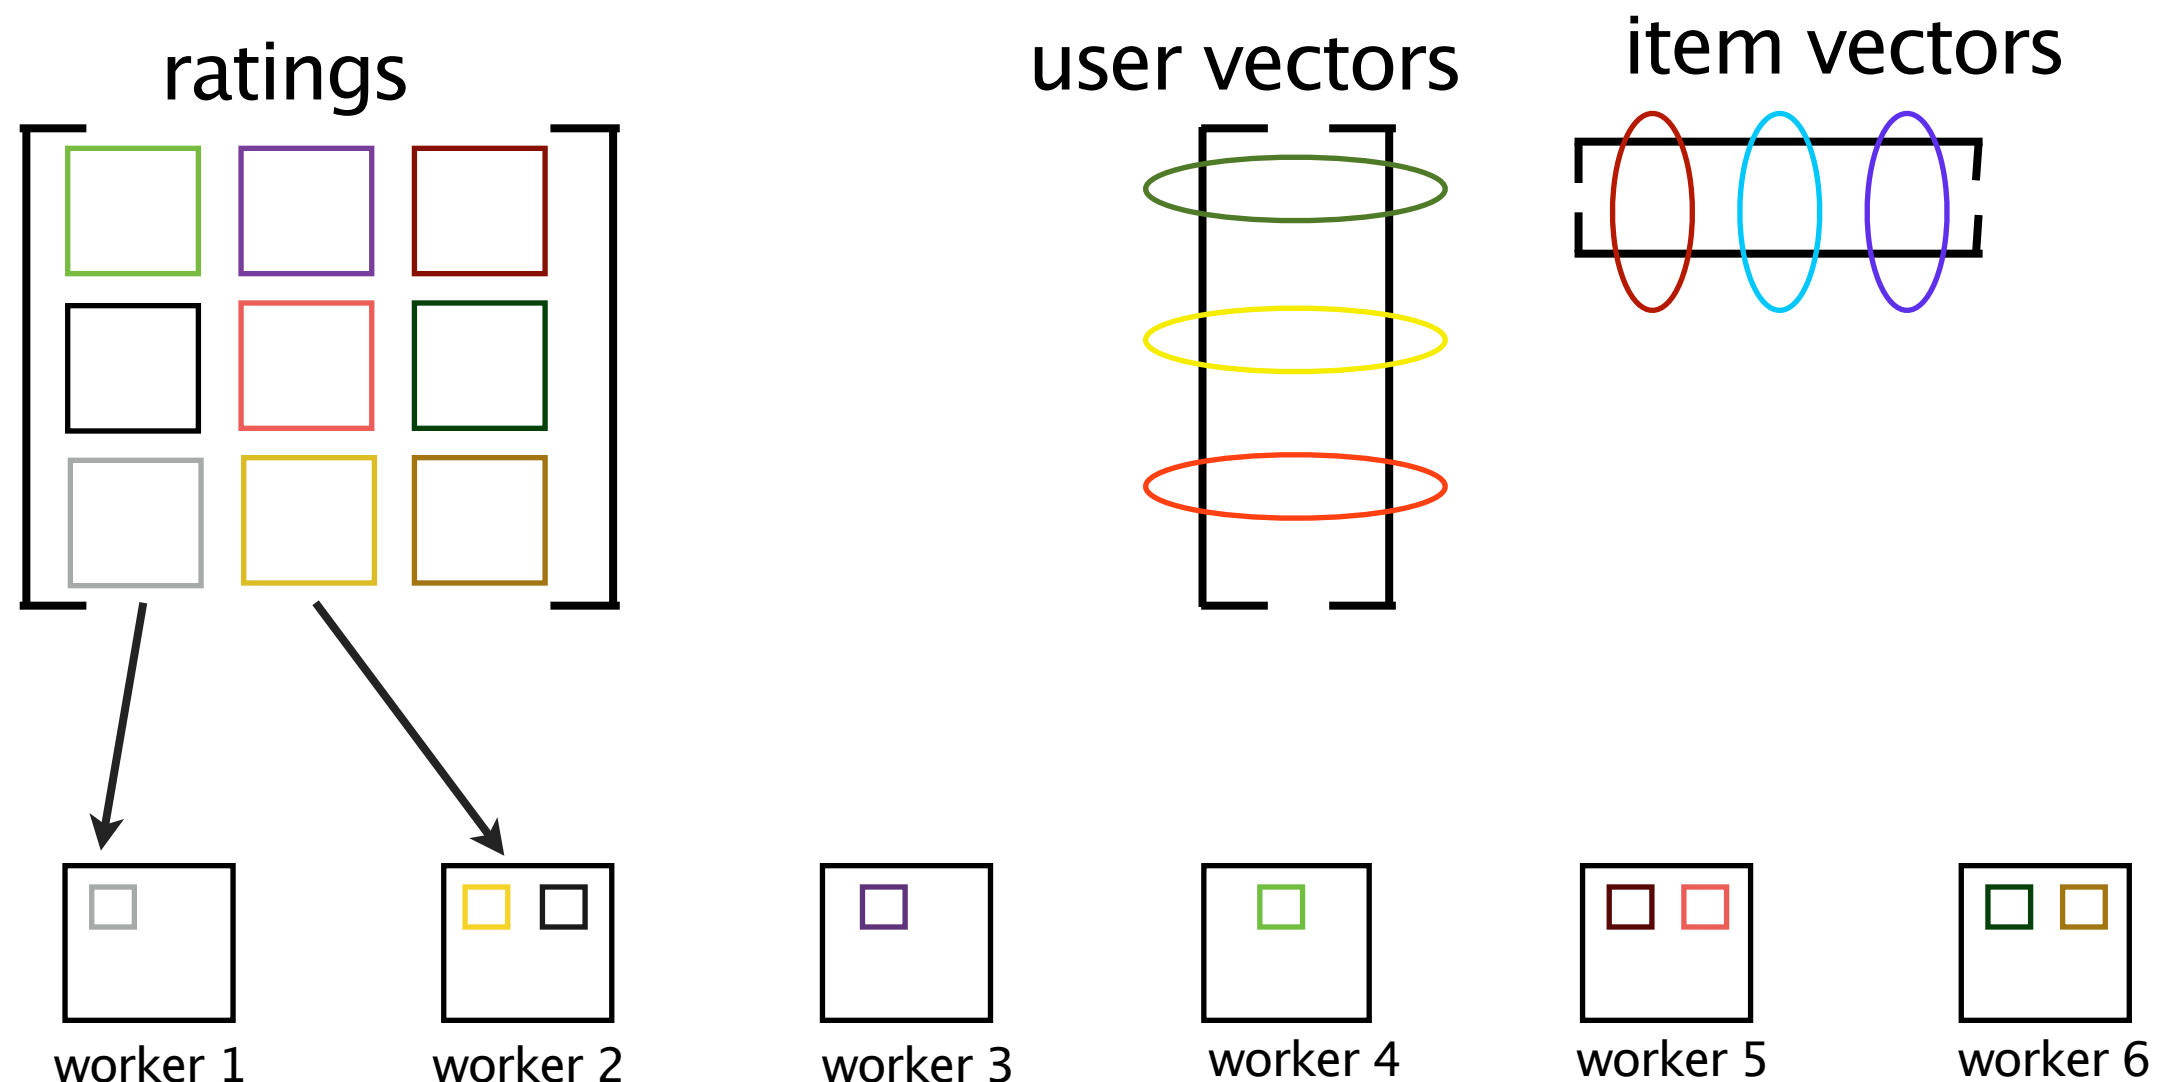
worker 5



worker 6

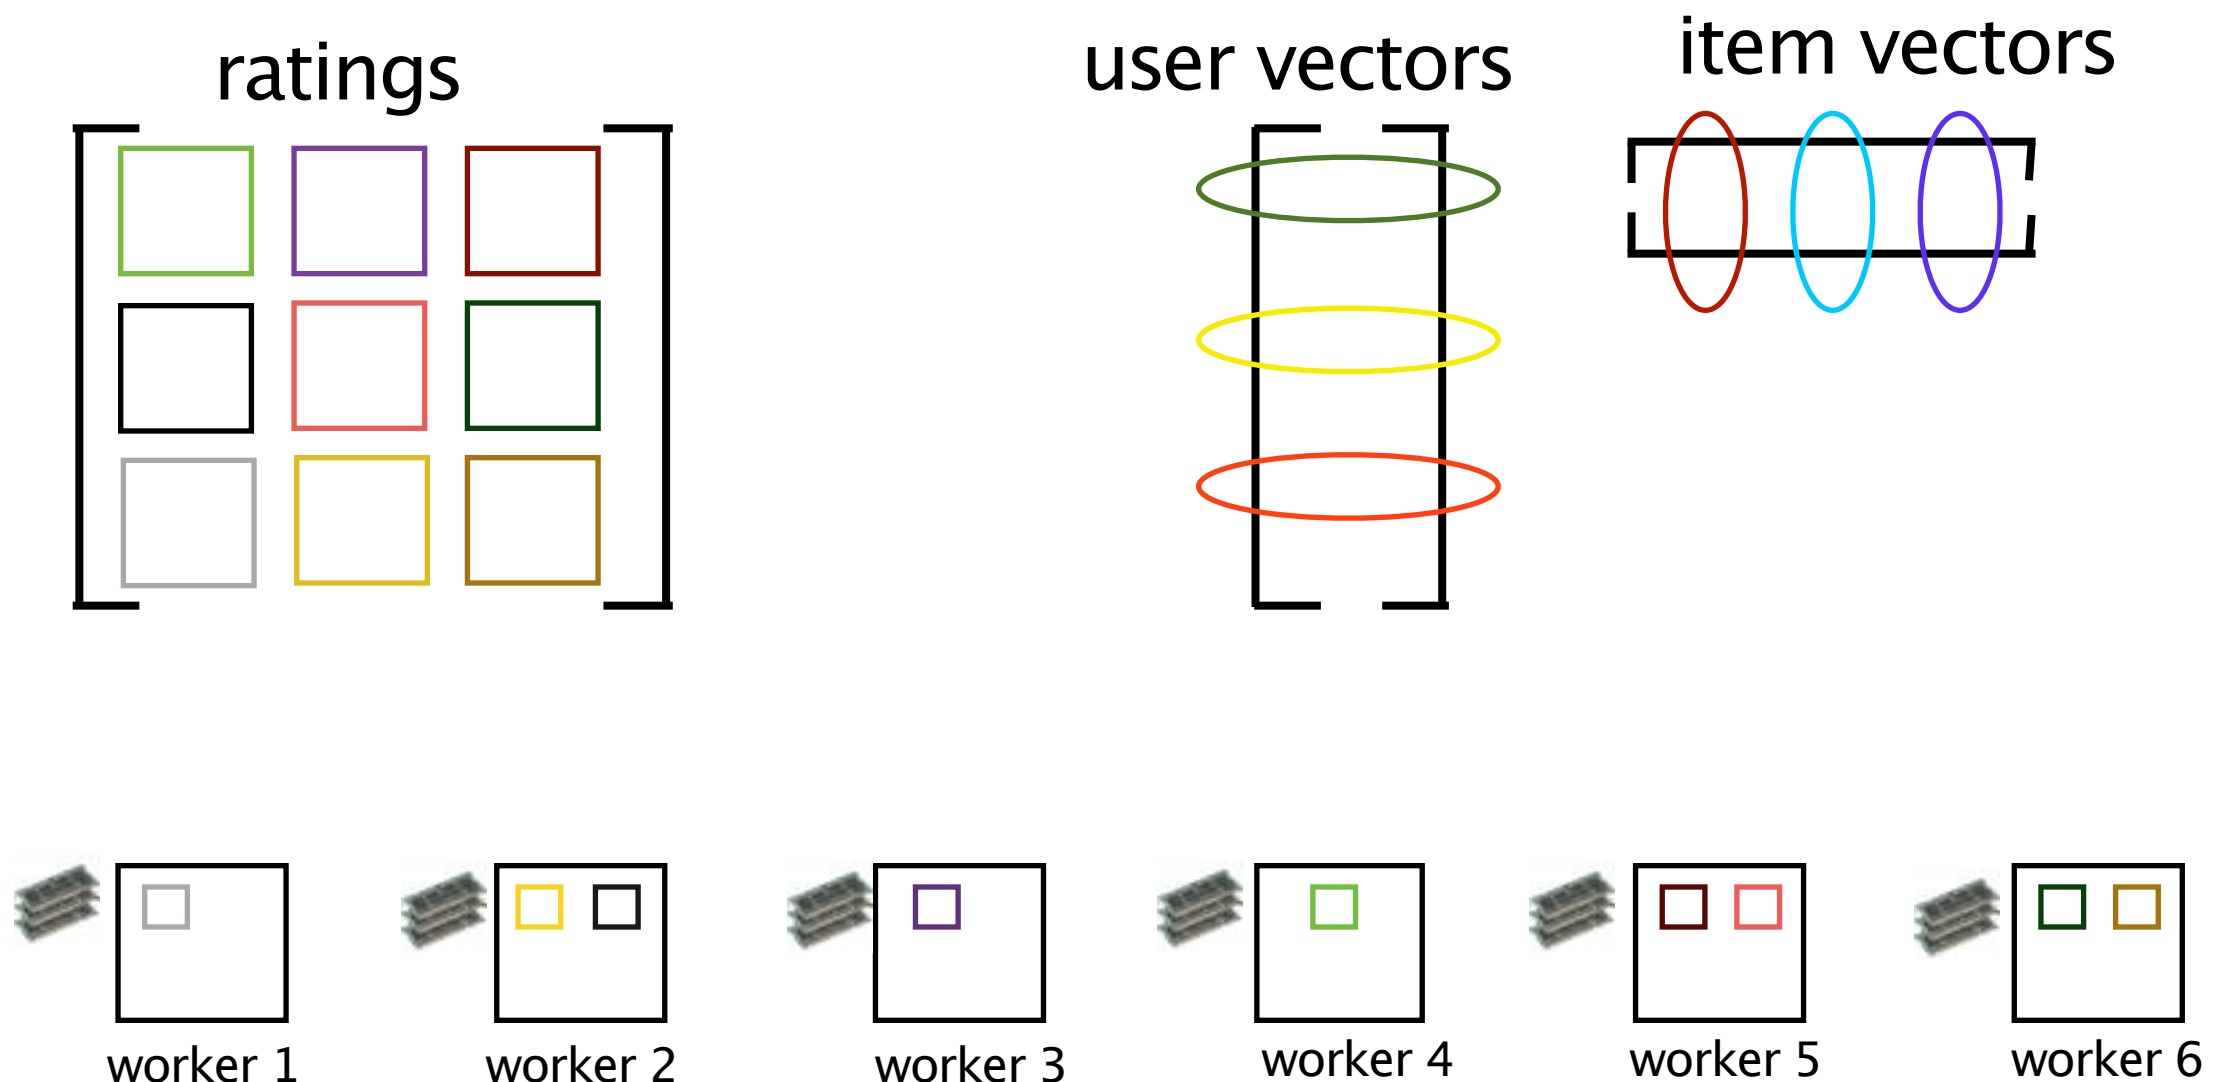
Second Attempt (full gridify)

- Group ratings matrix into $K \times L$, partition, and cache
- For each iteration:
 1. Compute YtY over item vectors and broadcast
 2. For each item vector send a copy to each rating block in the $\text{item} \% L$ column
 3. Compute intermediate terms for each block (partition)
 4. Group by user, aggregate intermediate terms, and solve for optimal user vector



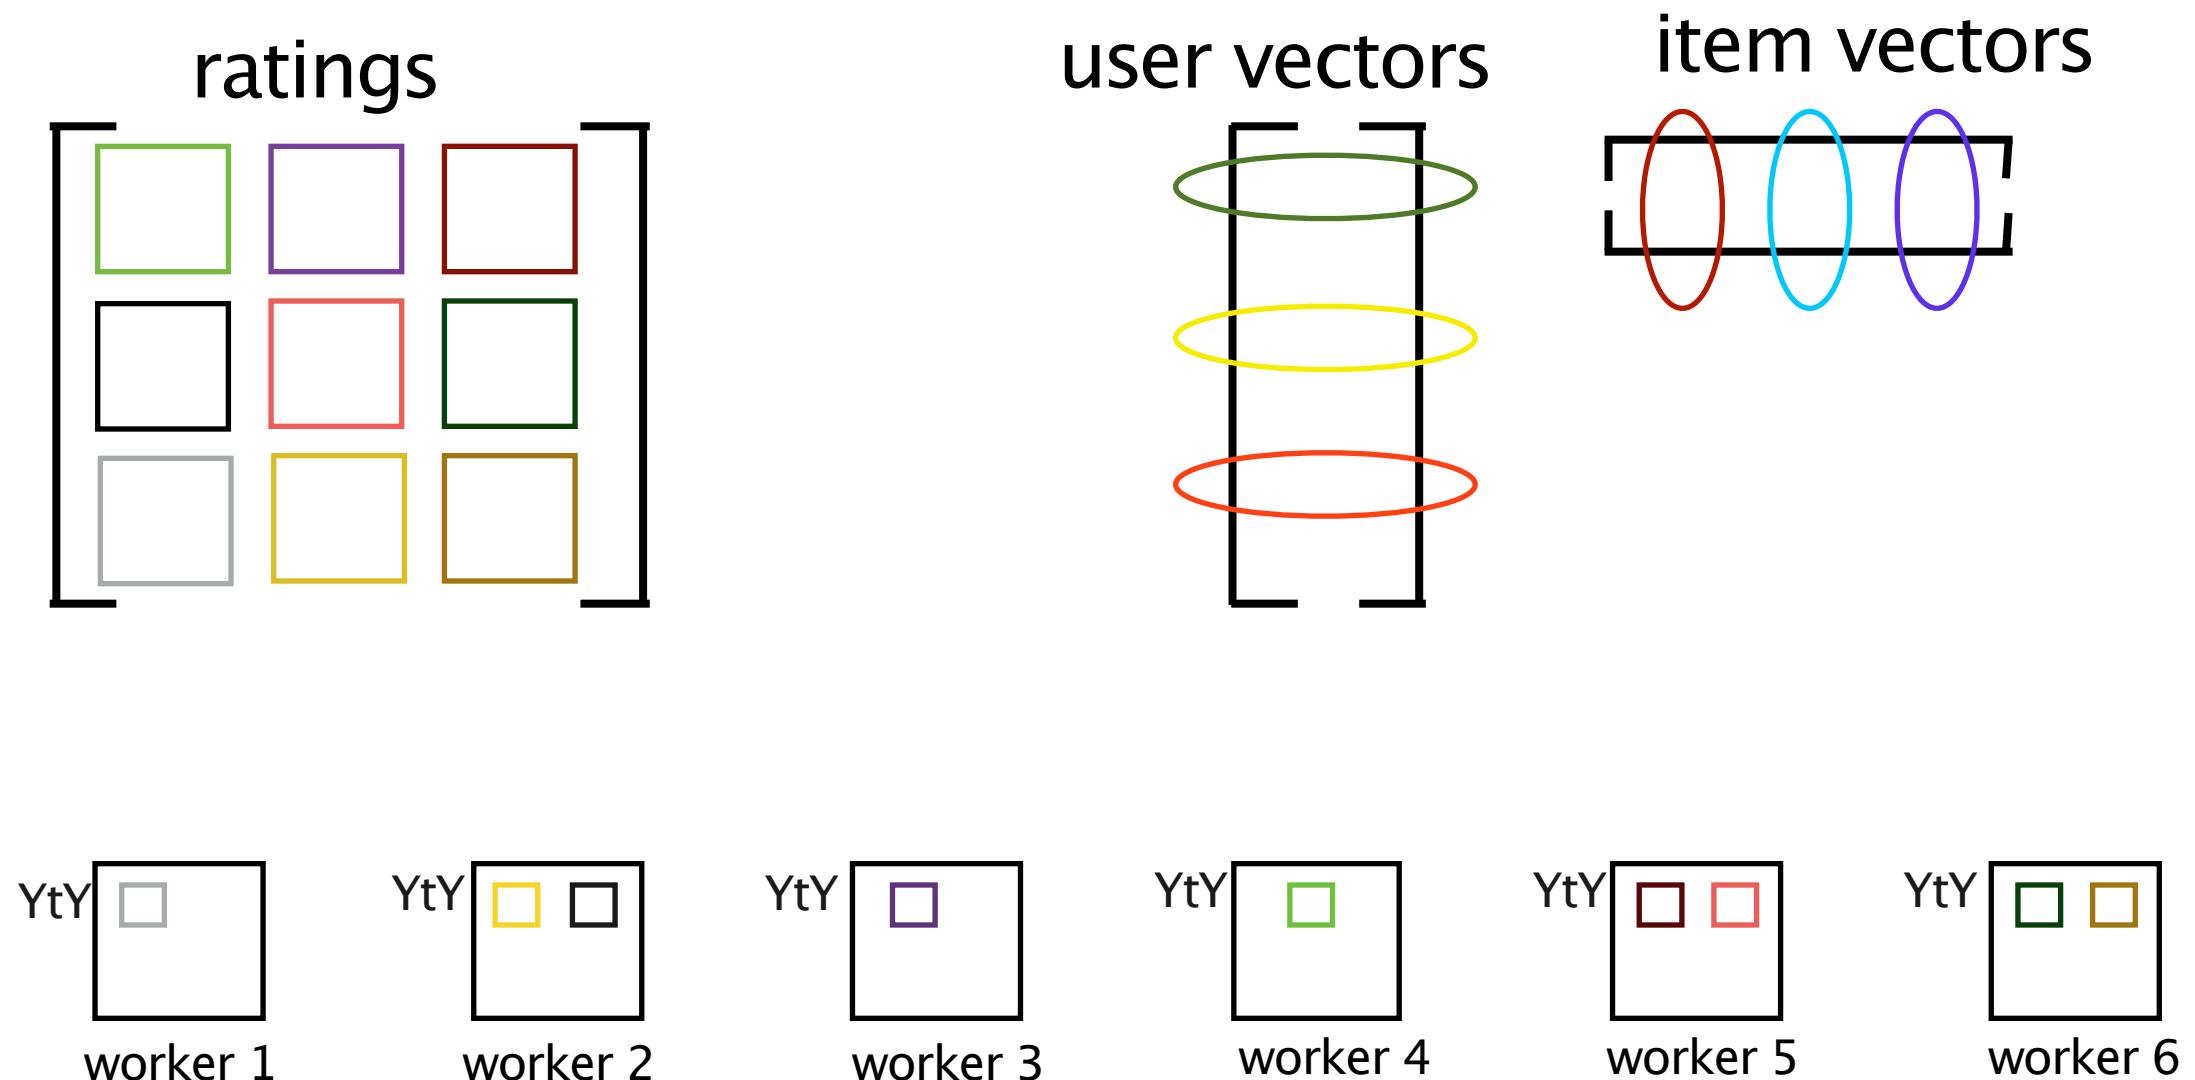
Second Attempt (full gridify)

- Group ratings matrix into $K \times L$, partition, and cache
- For each iteration:
 1. Compute YtY over item vectors and broadcast
 2. For each item vector send a copy to each rating block in the $\text{item} \% L$ column
 3. Compute intermediate terms for each block (partition)
 4. Group by user, aggregate intermediate terms, and solve for optimal user vector



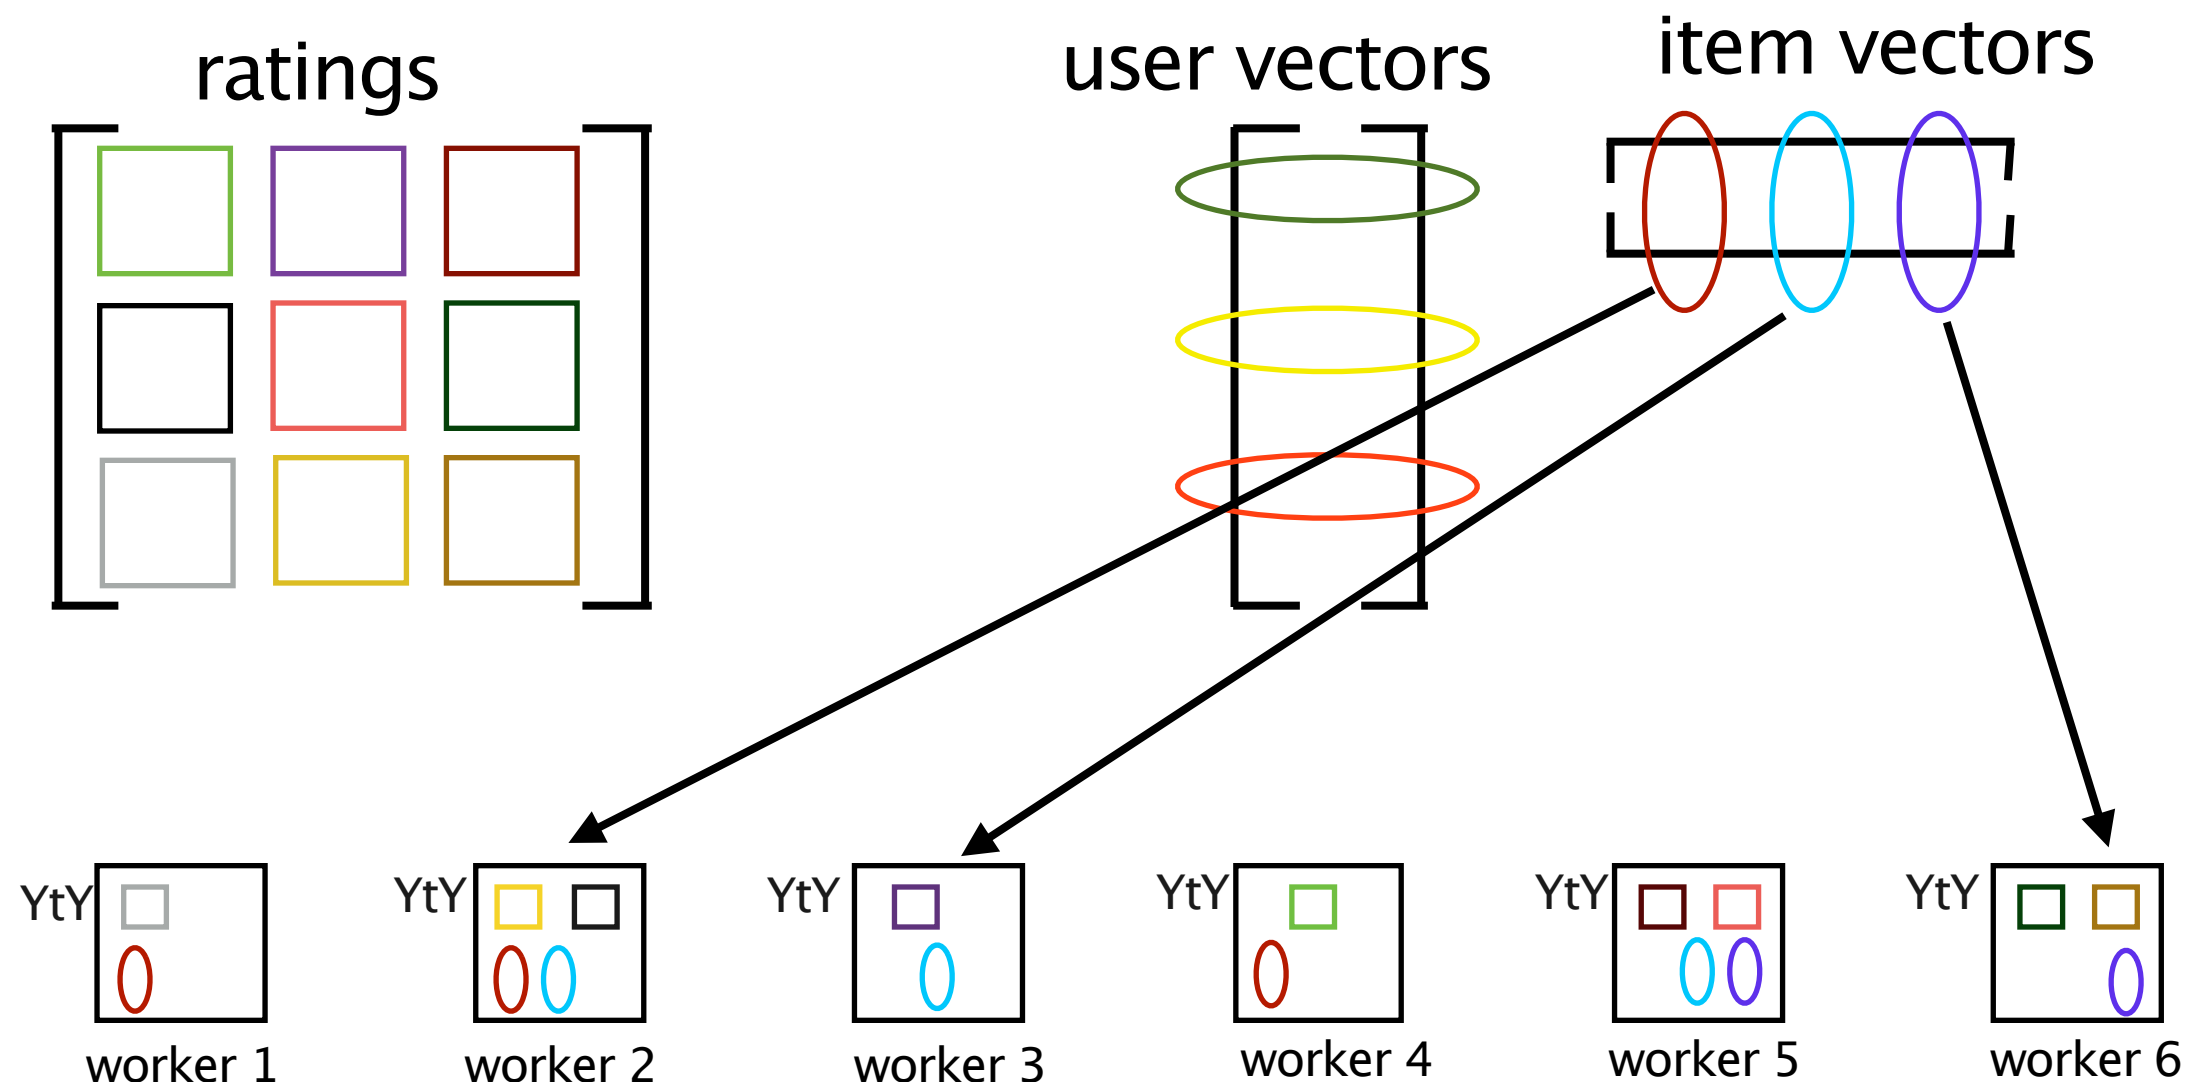
Second Attempt (full gridify)

- Group ratings matrix into $K \times L$, partition, and cache
- For each iteration:
 1. Compute $Y_t Y$ over item vectors and broadcast
 2. For each item vector send a copy to each rating block in the item % L column
 3. Compute intermediate terms for each block (partition)
 4. Group by user, aggregate intermediate terms, and solve for optimal user vector



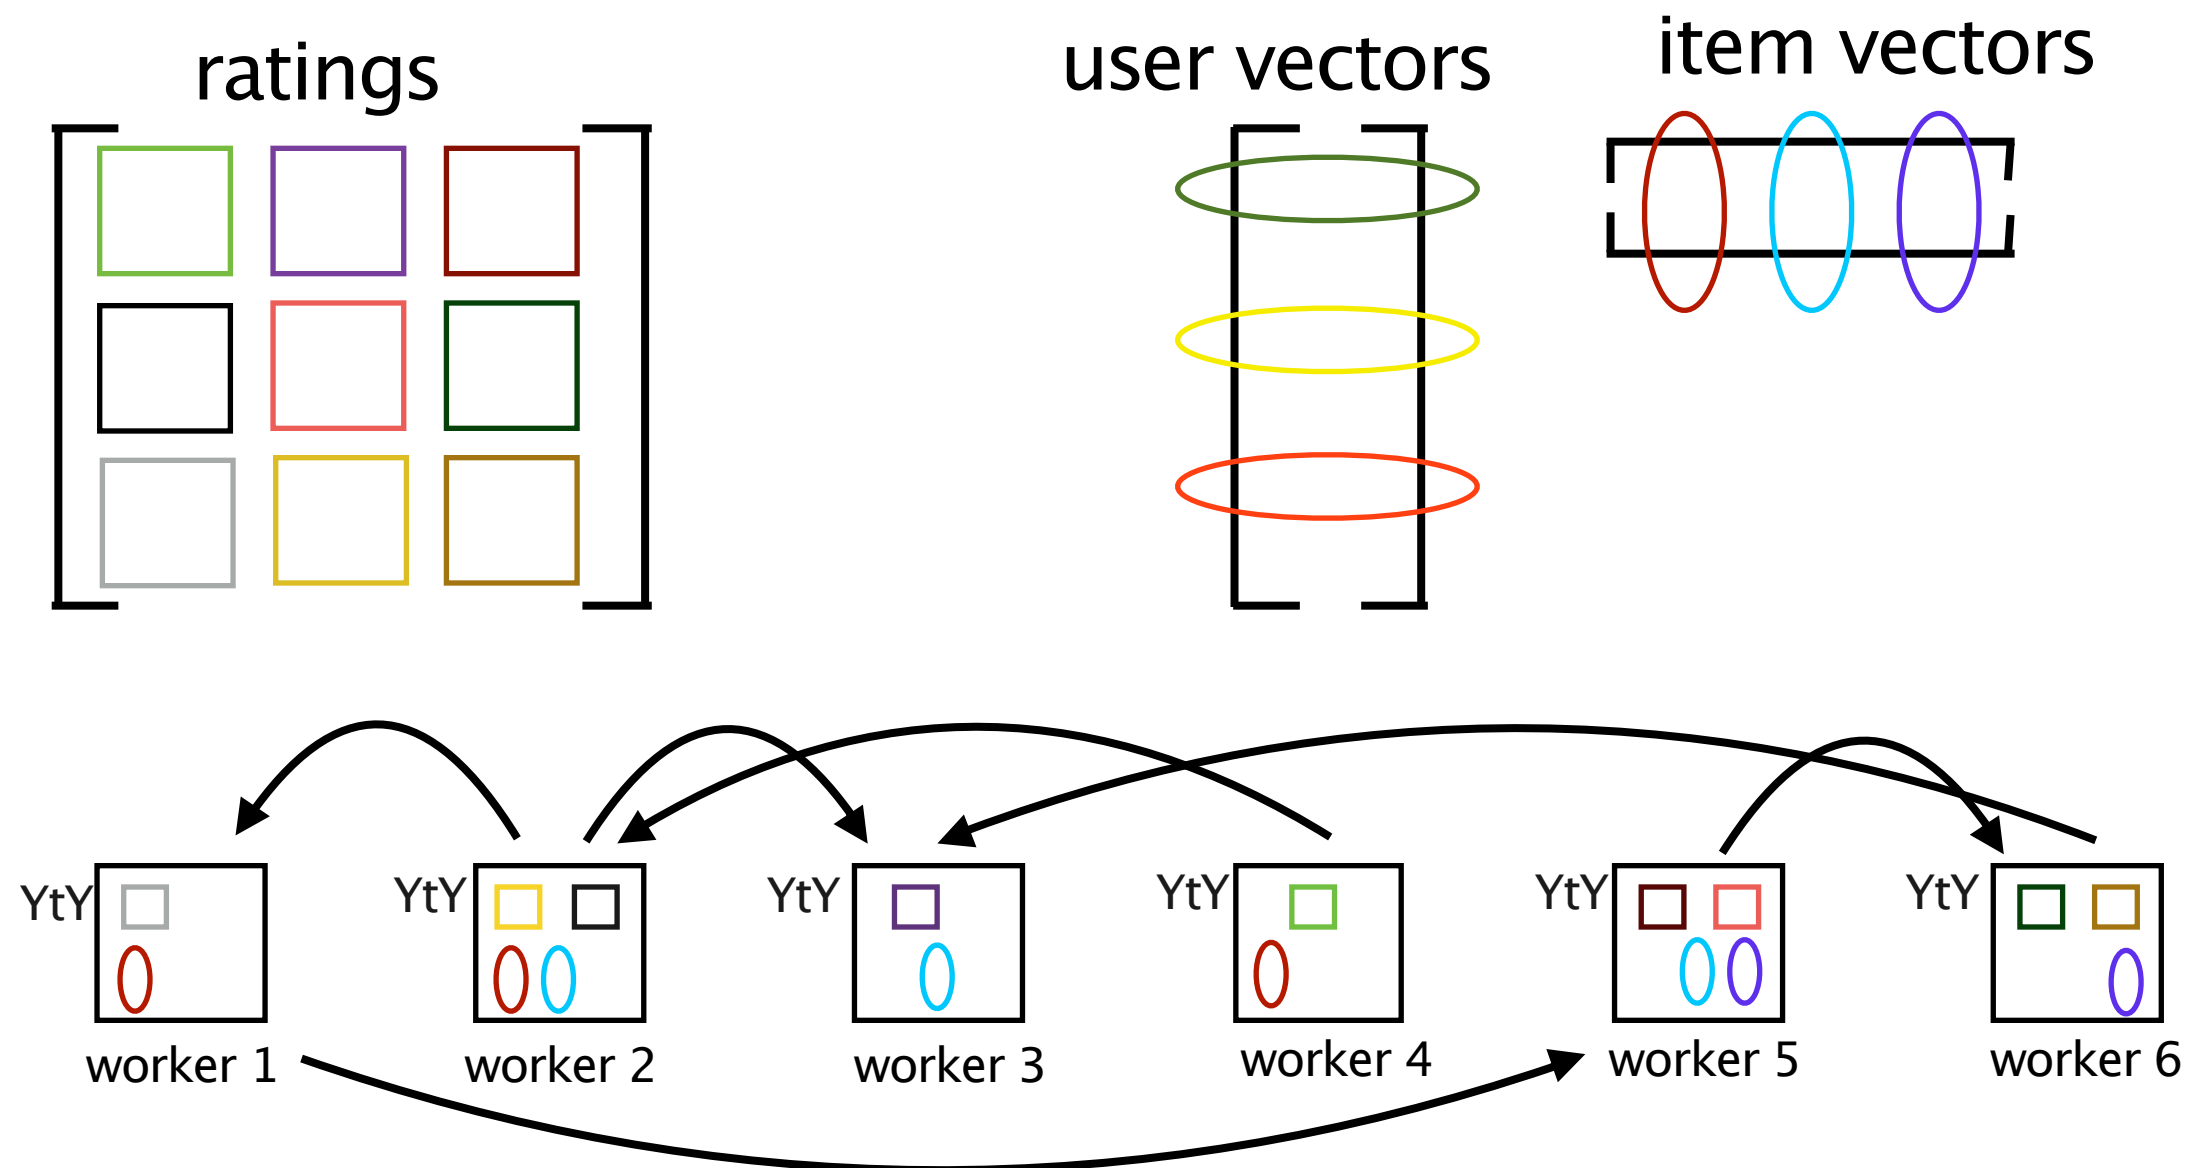
Second Attempt (full gridify)

- Group ratings matrix into $K \times L$, partition, and cache
- For each iteration:
 1. Compute $Y^t Y$ over item vectors and broadcast
 2. For each item vector send a copy to each rating block in the $\text{item} \% L$ column
 3. Compute intermediate terms for each block (partition)
 4. Group by user, aggregate intermediate terms, and solve for optimal user vector



Second Attempt (full gridify)

- Group ratings matrix into $K \times L$, partition, and cache
- For each iteration:
 1. Compute $Y^t Y$ over item vectors and broadcast
 2. For each item vector send a copy to each rating block in the $\text{item} \% L$ column
 3. Compute intermediate terms for each block (partition)
 4. Group by user, aggregate intermediate terms, and solve for optimal user vector




Second Attempt (full gridify)



- Group ratings matrix into $K \times L$, partition, and cache
- For each iteration:
 1. Compute YtY over item vectors and broadcast
 2. For each item vector send a copy to each rating block in the item % L column
 3. Compute intermediate terms for each block (partition)
 4. Group by user, aggregate intermediate terms, and solve for optimal user vector




user vectors




item vectors




worker 1




worker 2




worker 3



worker 4



worker 5



worker 6

Second Attempt

```
def fullGridify(ratings: RDD[Rating],
               k: Int,
               l: Int,
               partitioner: Partitioner) = {

  ratings
    .map{r: Rating =>
      val row = r.user % k
      val column = r.item % l
      (((row * l) + column), r)
    }.groupByKey(partitioner)
    .mapValues{itr: Iterable[Rating] =>
      itr.toList.groupBy(_.user)
    }.persist(StorageLevel.MEMORY_AND_DISK)
}

def updateVectors(ratingsByBlock: RDD[(Int, Map[Int, List[Rating]])],
                 itemsByBlock: RDD[(Int, Map[Int, VectorData])],
                 partitioner: Partitioner) = {

  val yty: DenseMatrix[Double] = computeYtY(itemsByBlock)

  val joinedVectorsRatings = joinVectorsRatings(ratingsByBlock, itemsByBlock, k, l, user, partitioner)

  val aggregatedTerms: RDD[(Int, Iterable[(Int, DenseMatrix[Double], DenseVector[Double])])] =
    aggregateTerms(joinedVectorsRatings, alpha, k, user, rank, partitioner)

  solveVectors(aggregatedTerms, yty, lambda, user)
}
```

Second Attempt

```
def fullGridify(ratings: RDD[Rating],
               k: Int,
               l: Int,
               partitioner: Partitioner) = {

  ratings
    .map{r: Rating =>
      val row = r.user % k
      val column = r.item % l
      (((row * l) + column), r)
    }.groupByKey(partitioner)
    .mapValues{itr: Iterable[Rating] =>
      itr.toList.groupBy(_.user)
    }.persist(StorageLevel.MEMORY_AND_DISK)
}

def updateVectors(ratingsByBlock: RDD[(Int, Map[Int, List[Rating]])],
                 itemsByBlock: RDD[(Int, Map[Int, VectorData])],
                 partitioner: Partitioner) = {

  val yty: DenseMatrix[Double] = computeYtY(itemsByBlock)

  val joinedVectorsRatings = joinVectorsRatings(ratingsByBlock, itemsByBlock, k, l, user, partitioner)

  val aggregatedTerms: RDD[(Int, Iterable[(Int, DenseMatrix[Double], DenseVector[Double])])] =
    aggregateTerms(joinedVectorsRatings, alpha, k, user, rank, partitioner)

  solveVectors(aggregatedTerms, yty, lambda, user)
}
```

● Pros


- Ratings get cached and never shuffled
- Each partition only requires a subset of item (or user) vectors in memory each iteration
- Potentially requires less local memory than a “half gridify” scheme

● Cons


- Sending lots of intermediate data over wire each iteration in order to aggregate and solve for optimal vectors
- More IO overhead than a “half gridify” scheme

Third Attempt (half gridify)


- Partition ratings matrix into K user (row) and item (column) blocks, partition, and cache
- For each iteration:
 1. Compute YtY over item vectors and broadcast
 2. For each item vector, send a copy to each user rating partition that requires it (potentially all partitions)
 3. Each partition aggregates intermediate terms and solves for optimal user vectors



user vectors




item vectors



worker 1




worker 2




worker 3



worker 4



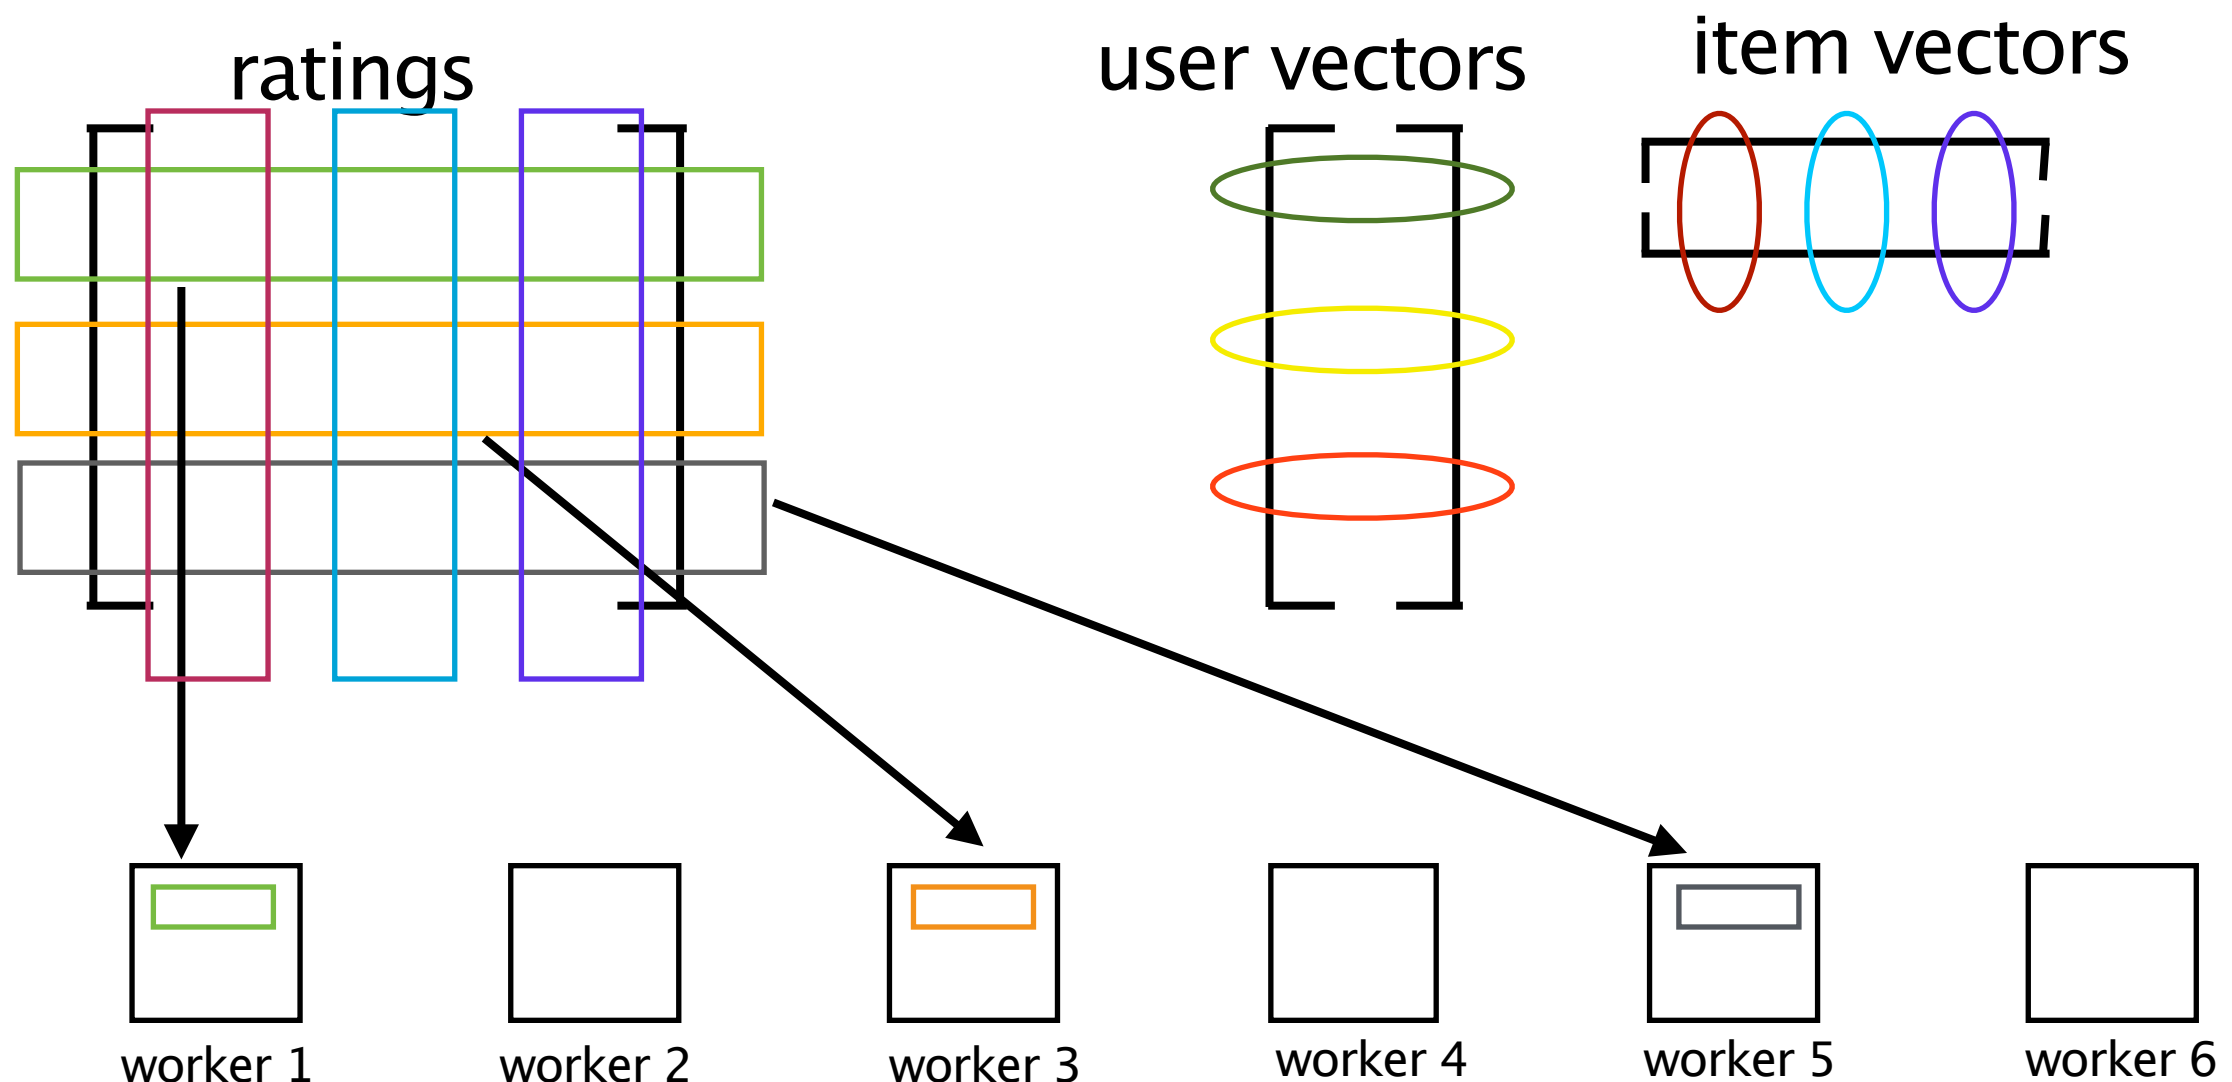
worker 5



worker 6


Third Attempt (half gridify)

- Partition ratings matrix into K user (row) and item (column) blocks, partition, and cache
- For each iteration:
 1. Compute YtY over item vectors and broadcast
 2. For each item vector, send a copy to each user rating partition that requires it (potentially all partitions)
 3. Each partition aggregates intermediate terms and solves for optimal user vectors




Third Attempt (half gridify)





- Partition ratings matrix into K user (row) and item (column) blocks, partition, and cache
- For each iteration:
 1. Compute YtY over item vectors and broadcast
 2. For each item vector, send a copy to each user rating partition that requires it (potentially all partitions)
 3. Each partition aggregates intermediate terms and solves for optimal user vectors



user vectors




item vectors




Third Attempt (half gridify)








- Partition ratings matrix into K user (row) and item (column) blocks, partition, and cache
- For each iteration:
 1. Compute $Y^t Y$ over item vectors and broadcast
 2. For each item vector, send a copy to each user rating partition that requires it (potentially all partitions)
 3. Each partition aggregates intermediate terms and solves for optimal user vectors



user vectors

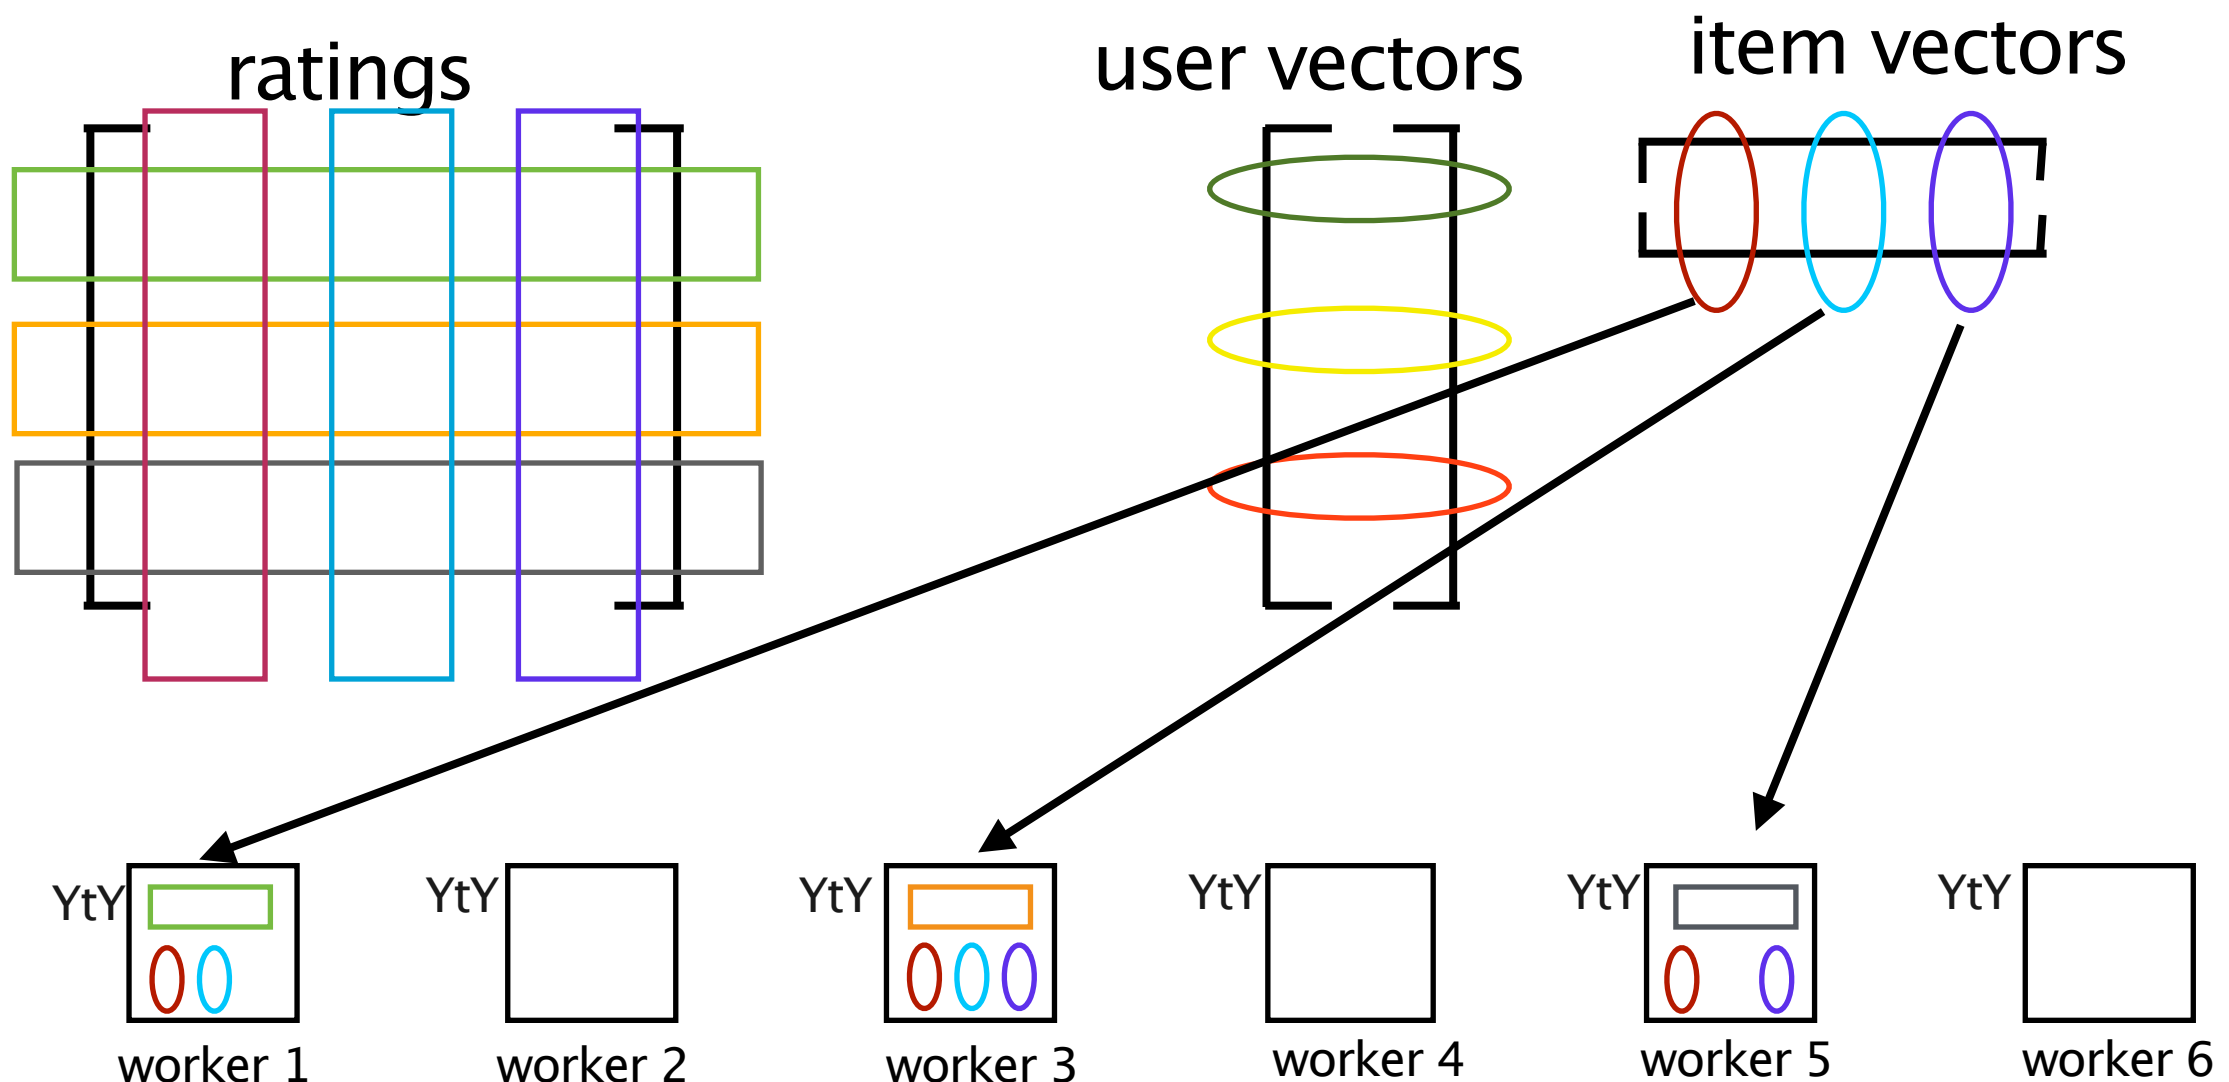


item vectors




Third Attempt (half gridify)

- Partition ratings matrix into K user (row) and item (column) blocks, partition, and cache
- For each iteration:
 1. Compute $Y^t Y$ over item vectors and broadcast
 2. For each item vector, send a copy to each user rating partition that requires it (potentially all partitions)
 3. Each partition aggregates intermediate terms and solves for optimal user vectors




Third Attempt (half gridify)








- Partition ratings matrix into K user (row) and item (column) blocks, partition, and cache
- For each iteration:
 1. Compute $Y^t Y$ over item vectors and broadcast
 2. For each item vector, send a copy to each user rating partition that requires it (potentially all partitions)
 3. Each partition aggregates intermediate terms and solves for optimal user vectors



user vectors




item vectors




Third Attempt (half gridify)


- Partition ratings matrix into K user (row) and item (column) blocks, partition, and cache
- For each iteration:
 1. Compute $Y^t Y$ over item vectors and broadcast
 2. For each item vector, send a copy to each user rating partition that requires it (potentially all partitions)
 3. Each partition aggregates intermediate terms and solves for optimal user vectors









user vectors



item vectors



Note that we removed the extra shuffle from the full gridify approach.



Third Attempt (half gridify)

```
private def updateFeatures(
  products: RDD[(Int, Array[Array[Double]])],
  productOutLinks: RDD[(Int, OutLinkBlock)],
  userInLinks: RDD[(Int, InLinkBlock)],
  partitioner: Partitioner,
  rank: Int,
  lambda: Double,
  alpha: Double,
  YtY: Option[Broadcast[DoubleMatrix]])
: RDD[(Int, Array[Array[Double]])] =
{
  val numBlocks = products.partitions.size
  productOutLinks.join(products).flatMap { case (bid, (outLinkBlock, factors)) =>
    val toSend = Array.fill(numBlocks)(new ArrayBuffer[Array[Double]])
    for (p <- 0 until outLinkBlock.elementIds.length; userBlock <- 0 until numBlocks) {
      if (outLinkBlock.shouldSend(p)(userBlock)) {
        toSend(userBlock) += factors(p)
      }
    }
    toSend.zipWithIndex.map { case (buf, idx) => (idx, (bid, buf.toArray)) }
  }.groupByKey(partitioner)
  .join(userInLinks)
  .mapValues { case (messages, inLinkBlock) =>
    updateBlock(messages, inLinkBlock, rank, lambda, alpha, YtY)
  }
}
```

Actual MLLib code!

- Pros
 - Ratings get cached and never shuffled
 - Once item vectors are joined with ratings partitions each partition has enough information to solve optimal user vectors without any additional shuffling/aggregation (which occurs with the “full gridify” scheme)
- Cons
 - Each partition could potentially require a copy of each item vector (which may not all fit in memory)
 - Potentially requires more local memory than “full gridify” scheme

ALS Running Times

- Dataset consisting of Spotify streaming data for 2 Million users and 500k artists
 - Note: full dataset consists of 40M users and 20M songs but we haven't yet successfully run with Spark
- All jobs run using 40 latent factors
- Spark jobs used 200 executors with 8G containers
- Hadoop job used 1k mappers, 300 reducers

Hadoop	Spark (full gridify)	Spark (half gridify)
10 hours	3.5 hours	1.5 hours

ALS Running Times

System	Wall-clock time (seconds)
MATLAB	15443
Mahout	4206
GraphLab	291
MLlib	481

- Dataset: scaled version of Netflix data (9X in size).
- Cluster: 9 machines.
- MLlib is an order of magnitude faster than Mahout.
- MLlib is within factor of 2 of GraphLab.



Fin













