

Information Retrieval

INFO 4300 / CS 4300

- Indexing
 - Inverted indexes
 - Compression
 - ➔ – Index construction
 - Ranking model

Index Construction

- Simple in-memory indexer


```
procedure BUILDINDEX(D)
  I ← HashTable()
  n ← 0
  for all documents d ∈ D do
    n ← n + 1
    T ← Parse(d)
    Remove duplicates from T
    for all tokens t ∈ T do
      if t ∉ I then
        It ← Array()
      end if
      It.append(n)
    end for
  end for
  return I
end procedure
```

▷ *D* is a set of text documents
▷ Inverted list storage
▷ Document numbering
▷ Parse document into tokens

Merging

- Merging addresses limited memory problem
 - Build the inverted list structure until memory runs out
 - Then write the partial index to disk, start making a new one
 - At the end of this process, the disk is filled with many partial indexes, which are merged
- Partial lists must be designed so they can be merged in small pieces
 - e.g., storing in alphabetical order

Merging



Distributed Indexing

- Distributed processing driven by need to index and analyze huge amounts of data (i.e., the Web)
- Large numbers of inexpensive servers used rather than larger, more expensive machines
- *MapReduce* is a distributed programming tool designed for indexing and analysis tasks

Example

- Given a large text file that contains data about credit card transactions
 - Each line of the file contains a credit card number and an amount of money
 - Determine the number of unique credit card numbers
- Could use hash table – memory problems
 - counting is simple with sorted file
- Similar with distributed approach
 - sorting and placement are crucial

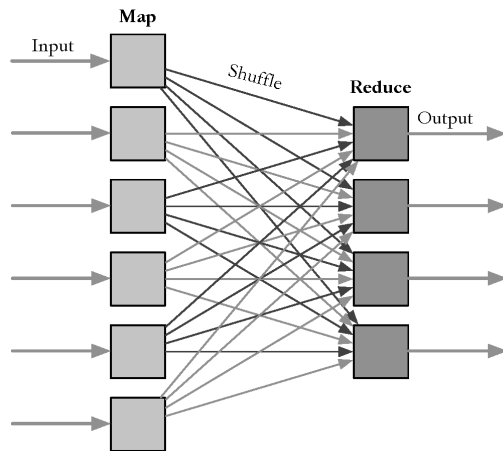
MapReduce

- Distributed programming framework that focuses on data placement and distribution
- *Mapper*
 - Generally, transforms a list of items into another list of items of the same length
- *Reducer*
 - Transforms a list of items into a single item
 - Definitions not so strict in terms of number of outputs
- Many mapper and reducer tasks on a cluster of machines

MapReduce

- Basic process
 - *Map* stage which transforms data records into pairs, each with a key and a value
 - *Shuffle* uses a hash function so that all pairs with the same key end up next to each other and on the same machine
 - *Reduce* stage processes records in batches, where all pairs with the same key are processed at the same time
- *Idempotence* of Mapper and Reducer provides fault tolerance
 - multiple operations on same input gives same output

MapReduce



Example

```
procedure MAPCREDITCARDS(input)
  while not input.done() do
    record ← input.next()
    card ← record.card
    amount ← record.amount
    Emit(card, amount)
  end while
end procedure

procedure REDUCECREDITCARDS(key, values)
  total ← 0
  card ← key
  while not values.done() do
    amount ← values.next()
    total ← total + amount
  end while
  Emit(card, total)
end procedure
```

Indexing Example

```
procedure MAPDOCUMENTSTOPOSTINGS(input)
  while not input.done() do
    document ← input.next()
    number ← document.number
    position ← 0
    tokens ← Parse(document)
    for each word w in tokens do
      Emit( number :position)
      position = position + 1
    end for
  end while
end procedure

procedure REDUCEPOSTINGSTOLISTS(key, values)
  word ← key
  WriteWord(word)
  while not input.done() do
    EncodePosting(values.next())
  end while
end procedure
```


Result Merging

- Index merging is a good strategy for handling updates when they come in large batches
- For small updates this is very inefficient
 - instead, create separate index for new documents, merge *results* from both searches
 - could be in-memory, fast to update and search
- Deletions handled using *delete list*
 - Modifications done by putting old version on delete list, adding new version to new documents index

Query Processing

- Document-at-a-time
 - Calculates complete scores for documents by processing all term lists, one document at a time
- Term-at-a-time
 - Accumulates scores for documents by processing term lists one at a time
- Both approaches have optimization techniques that significantly reduce time required to generate scores

Document-At-A-Time

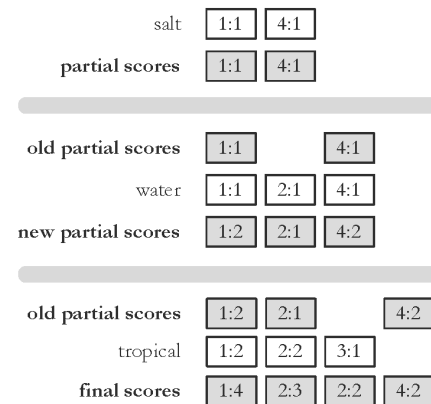


Document-At-A-Time

```

procedure DOCUMENTATATIMEREtrieval( $Q, I, f, g, k$ )
   $L \leftarrow \text{Array}()$ 
   $R \leftarrow \text{PriorityQueue}(k)$ 
  for all terms  $w_i$  in  $Q$  do
     $l_i \leftarrow \text{InvertedList}(w_i, I)$ 
     $L.\text{add}(l_i)$ 
  end for
  for all documents  $d \in I$  do
    for all inverted lists  $l_i$  in  $L$  do
      if  $l_i$  points to  $d$  then
         $s_D \leftarrow s_D + g_i(Q)f_i(l_i)$        $\triangleright$  Update the document score
         $l_i.\text{movePastDocument}(d)$ 
      end if
    end for
     $R.\text{add}(s_D, D)$ 
  end for
  return the top  $k$  results from  $R$ 
end procedure
    
```

Term-At-A-Time



Term-At-A-Time

```

procedure TERMATATIMERETRIEVAL( $Q, I, f, g, k$ )
   $A \leftarrow$  HashTable()
   $L \leftarrow$  Array()
   $R \leftarrow$  PriorityQueue( $k$ )
  for all terms  $w_i$  in  $Q$  do
     $l_i \leftarrow$  InvertedList( $w_i, I$ )
     $L.add( l_i )$ 
  end for
  for all lists  $l_i \in L$  do
    while  $l_i$  is not finished do
       $d \leftarrow l_i.get\text{CurrentDocument}()$ 
       $A_d \leftarrow A_d + g_i(Q)f(l_i)$ 
       $l_i.moveTo\text{NextDocument}()$ 
    end while
  end for
  for all accumulators  $A_d$  in  $A$  do
     $s_D \leftarrow A_d$  ▷ Accumulator contains the document score
     $R.add( s_D, D )$ 
  end for
  return the top  $k$  results from  $R$ 
end procedure

```

Optimization Techniques

- Term-at-a-time uses more memory for accumulators, but accesses disk more efficiently
- Two classes of optimization
 - Read less data from inverted lists
 - » e.g., skip lists
 - » better for simple feature functions
 - Calculate scores for fewer documents
 - » e.g., conjunctive processing
 - » better for complex feature functions

```

1: procedure TERMATATIMERETRIEVAL( $Q, I, f, g, k$ )
2:    $A \leftarrow$  HashTable()
3:    $L \leftarrow$  Array()
4:    $R \leftarrow$  PriorityQueue( $k$ )
5:   for all terms  $w_i$  in  $Q$  do
6:      $l_i \leftarrow$  InvertedList( $w_i, I$ )
7:      $L.add( l_i )$ 
8:   end for
9:   for all lists  $l_i \in L$  do
10:    while  $l_i$  is not finished do
11:     if  $i = 0$  then
12:        $d \leftarrow l_i.get\text{CurrentDocument}()$ 
13:        $A_d \leftarrow A_d + g_i(Q)f(l_i)$ 
14:     else
15:        $d \leftarrow l_i.get\text{CurrentDocument}()$ 
16:        $d \leftarrow A_d.get\text{NextDocumentAfter}(d)$ 
17:        $l_i.skip\text{ForwardTo}(d)$ 
18:       if  $l_i.get\text{CurrentDocument}() = d$  then
19:          $A_d \leftarrow A_d + g_i(Q)f(l_i)$ 
20:       else
21:          $A.remove(d)$ 
22:       end if
23:     end if
24:   end while
25: end for
26: for all accumulators  $A_d$  in  $A$  do
27:    $s_D \leftarrow A_d$  ▷ Accumulator contains the document score
28:    $R.add( s_D, D )$ 
29: end for
30: return the top  $k$  results from  $R$ 
31: end procedure

```

Conjunctive Term-at-a-Time

```

1: procedure DOCUMENTATATIMERETRIEVAL( $Q, I, f, g, k$ )
2:    $L \leftarrow$  Array()
3:    $R \leftarrow$  PriorityQueue( $k$ )
4:   for all terms  $w_i$  in  $Q$  do
5:      $l_i \leftarrow$  InvertedList( $w_i, I$ )
6:      $L.add( l_i )$ 
7:   end for
8:   while all lists in  $L$  are not finished do
9:     for all inverted lists  $l_i$  in  $L$  do
10:      if  $l_i.get\text{CurrentDocument}() > d$  then
11:         $d \leftarrow l_i.get\text{CurrentDocument}()$ 
12:      end if
13:     end for
14:     for all inverted lists  $l_i$  in  $L$  do  $l_i.skip\text{ForwardToDocument}(d)$ 
15:     if  $l_i$  points to  $d$  then
16:        $s_d \leftarrow s_d + g_i(Q)f_i(l_i)$  ▷ Update the document score
17:        $l_i.move\text{PastDocument}( d )$ 
18:     else
19:       break
20:     end if
21:   end for
22:    $R.add( s_d, d )$ 
23: end while
24: return the top  $k$  results from  $R$ 
25: end procedure

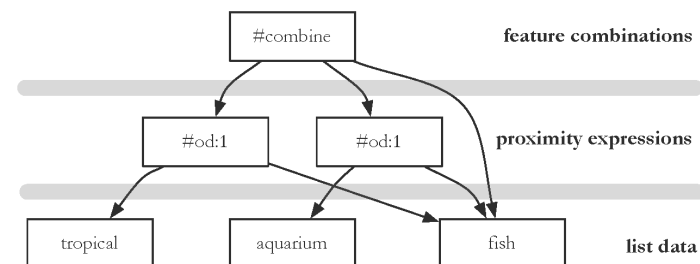
```

Conjunctive Document-at-a-Time

Structured Queries

- *Query language* can support specification of complex features
 - similar to SQL for database systems
 - *query translator* converts the user's input into the structured query representation
 - Galago query language is the example used here
 - e.g., Galago query:
`#combine(#od:1(tropical fish) #od:1(aquarium fish) fish)`

Evaluation Tree for Structured Query



Distributed Evaluation

- Basic process
 - All queries sent to a *director machine*
 - Director then sends messages to many *index servers*
 - Each index server does some portion of the query processing
 - Director organizes the results and returns them to the user
- Two main approaches
 - Document distribution
 - » by far the most popular
 - Term distribution

Distributed Evaluation

- Document distribution
 - each index server acts as a search engine for a small fraction of the total collection
 - director sends a copy of the query to each of the index servers, each of which returns the top-*k* results
 - results are merged into a single ranked list by the director
- Collection statistics should be shared for effective ranking

Distributed Evaluation

- Term distribution
 - Single index is built for the whole cluster of machines
 - Each inverted list in that index is then assigned to one index server
 - » in most cases the data to process a query is not stored on a single machine
 - One of the index servers is chosen to process the query
 - » usually the one holding the longest inverted list
 - Other index servers send information to that server
 - Final results sent to director

Caching

- Query distributions similar to Zipf
 - About ½ each day are unique, but some are very popular
- Caching can significantly improve effectiveness
 - Cache popular query results
 - Cache common inverted lists
- Inverted list caching can help with unique queries
- Cache must be refreshed to prevent stale data