

Pig Cookbook

Table of contents

1 Overview.....	2
2 Performance Enhancers.....	2

1. Overview

This document provides hints and tips for pig users.

2. Performance Enhancers

2.1. Use Optimization

Pig supports various [optimization rules](#) which are turned on by default. Become familiar with these rules.

2.2. Use Types

If types are not specified in the load statement, Pig assumes the type of =double= for numeric computations. A lot of the time, your data would be much smaller, maybe, integer or long. Specifying the real type will help with speed of arithmetic computation. It has an additional advantage of early error detection.

```
--Query 1
A = load 'myfile' as (t, u, v);
B = foreach A generate t + u;

--Query 2
A = load 'myfile' as (t: int, u: int, v);
B = foreach A generate t + u;
```

The second query will run more efficiently than the first. In some of our queries with see 2x speedup.

2.3. Project Early and Often

Pig does not (yet) determine when a field is no longer needed and drop the field from the row. For example, say you have a query like:

```
A = load 'myfile' as (t, u, v);
B = load 'myotherfile' as (x, y, z);
C = join A by t, B by x;
D = group C by u;
E = foreach D generate group, COUNT($1);
```

There is no need for v, y, or z to participate in this query. And there is no need to carry both t and x past the join, just one will suffice. Changing the above query to

```
A = load 'myfile' as (t, u, v);
```

```
A1 = foreach A generate t, u;  
B = load 'myotherfile' as (x, y, z);  
B1 = foreach B generate x;  
C = join A1 by t, B1 by x;  
C1 = foreach C generate t, u;  
D = group C1 by u;  
E = foreach D generate group, COUNT($1);
```

will greatly reduce the amount of data being carried through the map and reduce phases by pig. Depending on your data, this can produce significant time savings. In queries similar to the example given we have seen total time drop by 50%.

2.4. Filter Early and Often

As with early projection, in most cases it is beneficial to apply filters as early as possible to reduce the amount of data flowing through the pipeline.

```
-- Query 1  
A = load 'myfile' as (t, u, v);  
B = load 'myotherfile' as (x, y, z);  
C = filter A by t == 1;  
D = join C by t, B by x;  
E = group D by u;  
F = foreach E generate group, COUNT($1);  
  
-- Query 2  
A = load 'myfile' as (t, u, v);  
B = load 'myotherfile' as (x, y, z);  
C = join A by t, B by x;  
D = group C by u;  
E = foreach D generate group, COUNT($1);  
F = filter E by C.t == 1;
```

The first query is clearly more efficient than the second one because it reduces the amount of data going into the join.

One case where pushing filters up might not be a good idea is if the cost of applying filter is very high and only a small amount of data is filtered out.

2.5. Reduce Your Operator Pipeline

For clarity of your script, you might choose to split your projects into several steps for instance:

```
A = load 'data' as (in: map[]);  
-- get key out of the map  
B = foreach A generate in#k1 as k1, in#k2 as k2;  
-- concatenate the keys  
C = foreach B generate CONCAT(k1, k2);
```

```
.....
```

While the example above is easier to read, you might want to consider combining the two foreach statements to improve your query performance:

```
A = load 'data' as (in: map[]);
-- concatenate the keys from the map
B = foreach A generate CONCAT(in#k1, in#k2);
....
```

The same goes for filters.

2.6. Make your UDFs Algebraic

Queries that can take advantage of the combiner generally ran much faster (sometimes several times faster) than the versions that don't. The latest code significantly improves combiner usage; however, you need to make sure you do your part. If you have a UDF that works on grouped data and is, by nature, algebraic (meaning their computation can be decomposed into multiple steps) make sure you implement it as such. For details on how to write algebraic UDFs, see the [Pig UDF Manual](#).

```
A = load 'data' as (x, y, z)
B = group A by x;
C = foreach B generate group, MyUDF(A);
....
```

If MyUDF is algebraic, the query will use combiner and run much faster. You can run `explain` command on your query to make sure that combiner is used.

2.7. Drop Nulls Before a Join

This comment only applies to pig 0.2.0 branch, as pig 0.1.0 does not have nulls.

With the introduction of nulls, join and cogroup semantics were altered to work with nulls. The semantic for cogrouping with nulls is that nulls from a given input are grouped together, but nulls across inputs are not grouped together. This preserves the semantics of grouping (nulls are collected together from a single input to be passed to aggregate functions like COUNT) and the semantics of join (nulls are not joined across inputs). Since flattening an empty bag results in an empty row, in a standard join the rows with a null key will always be dropped. The join:

```
A = load 'myfile' as (t, u, v);
B = load 'myotherfile' as (x, y, z);
C = join A by t, B by x;
```

is rewritten by pig to

```
A = load 'myfile' as (t, u, v);
B = load 'myotherfile' as (x, y, z);
C1 = cogroup A by t INNER, B by x INNER;
C = foreach C1 generate flatten(A), flatten(B);
```

Since the nulls from A and B won't be collected together, when the nulls are flattened we're guaranteed to have an empty bag, which will result in no output. So the null keys will be dropped. But they will not be dropped until the last possible moment. If the query is rewritten to

```
A = load 'myfile' as (t, u, v);
B = load 'myotherfile' as (x, y, z);
A1 = filter A by t is not null;
B1 = filter B by x is not null;
C = join A1 by t, B1 by x;
```

then the nulls will be dropped before the join. Since all null keys go to a single reducer, if your key is null even a small percentage of the time the gain can be significant. In one test where the key was null 7% of the time and the data was spread across 200 reducers, we saw a about a 10x speed up in the query by adding the early filters.

2.8. Take Advantage of Join Optimization

The optimization insures that the last table in the join is not brought into memory but stream through instead. The optimization reduces the amount of memory used which means you can avoid spilling the data and also should be able to scale your query to larger data volumes.

To take advantage of this optimization, make sure that the table with the largest number of tuples per key is the last table in your query.

```
small = load 'small_file' as (t, u, v);
large = load 'large_file' as (x, y, z);
C = join small by t, large by x;
```

In some of our tests we saw 10x performance improvement as the result of this optimization.

2.9. Use Fragment Replicate Join

This type of join works well if one of tables is small enough to fit into main memory. In this case, Pig can perform a very efficient join since, in hadoop world, it can be done completely on the map side.

```
tiny = load 'small_file' as (t, u, v);
large = load 'large_file' as (x, y, z);
C = join large by x, tiny by t using "replicated";
```

Note that the large table must come first followed by one or more small tables. All small

tables together must fit into main memory.

This feature is new and experimental. It is experimental because we don't have a strong sense of how small the small table must be to fit into memory. In our tests with a simple query that involved just join a table of up to 100M can be used if the process overall gets 1 GB of memory. If the table does not fit into memory, the process would fail and generate an error.

2.10. Use PARALLEL Keyword

PARALLEL controls the number of reducers invoked by Hadoop. The default out of the box is 1 which, in most cases, is not what you want. A reasonable heuristic to use is something like

```
<num machines> * <num reduce slots per machine> * 0.9
```

The keyword makes sense on any operator that starts a reduce phase. This includes GROUP, COGROUP, JOIN, DISTINCT, LIMIT, ORDER BY.

Example:

```
A = load 'myfile' as (t, u, v);
B = group A by t PARALLEL 18;
.....
```

2.11. Use LIMIT

A lot of the times, you are not interested in the entire output but either a sample or top results. In those cases, using LIMIT can yield a much better performance as we push the limit as high as possible to minimize the amount of data travelling through the pipeline.

Sample:

```
A = load 'myfile' as (t, u, v);
B = limit A 500;
```

Top results:

```
A = load 'myfile' as (t, u, v);
B = order A by t;
C = limit B 500;
```

2.12. Prefer DISTINCT over GROUP BY - GENERATE

When it comes to extracting the unique values from a column in a relation, one of two approaches can be used:

Example Using GROUP BY - GENERATE

```
A = load 'myfile' as (t, u, v);
B = foreach A generate u;
C = group B by u;
D = foreach C generate group as uniquekey;
dump D;
```

Example Using DISTINCT

```
A = load 'myfile' as (t, u, v);
B = foreach A generate u;
C = distinct B;
dump C;
```

In pig 0.1.x, DISTINCT is just GROUP BY/PROJECT under the hood. In pig 0.2.0 it is not, and it is much faster and more efficient (depending on your key cardinality, up to 20x faster in pig team's tests). Therefore, the use of DISTINCT is recommended over GROUP BY - GENERATE.