

The SQL Optimizer Project: _Method and _Tree in SAS®9.1

Russ Lavery
(Thanks to Paul Sherman)

ABSTRACT

This paper discusses two little known options of SAS® Proc SQL: _method and _tree. The main body of information, and the major opportunity to learn about the topics, exists in the heavily annotated appendix. Efficient use of this material might involve reading, and printing, this paper then working through the appendix. This paper will also attempt to describe some of the logic that the Optimizer employs.

Proc SQL has a powerful subroutine, the SQL Optimizer, that examines submitted SQL code and the state of the system (file size, index presence, buffersize, sort order etc). The SQL Optimizer creates a “run plan” for optimally running the query. Run plans describe executable programs that the Optimizer will create to produce the desired output. These executable programs can be quite complicated and often involve the creating, sorting and merging of many temporary files. Consistent with the Optimizer’s goal of minimizing run times, the executable programs will trim variables and observations from the input file(s)/working file(s) as soon as they can be removed.

Many details of the run plan can be determined by using two Proc SQL options (_Method and _Tree) and this paper explains output from these two options. _Method and _Tree produce different output that present different aspects of the run plan. Learning to interpret _Method and _Tree can help programmers explain why small variations in code, or system conditions, can cause substantial variations in run time.

INTRODUCTION

This paper introduces two little known options of SAS Proc SQL: _Method and _Tree. The majority of the information, and the opportunity to learn about the topics, exists in the heavily annotated appendix. Reading the appendix is strongly recommended.

In SQL, you tell SQL what you want for results, not how you want the results produced. SAS SQL has a powerful subroutine, the SQL Optimizer, that decides how the SQL query should be executed in order to minimize run time. The Optimizer examines submitted SQL code and characteristics of the SAS system and then creates efficient executable statements for the submitted query. The created code can be quite complicated and often involves the creating, sorting and merging of many temporary files as well as the trimming of variables and observation at times that will minimize run time.

All versions of SQL have Optimizers and, while they perform very well, performance can sometimes be improved by manual intervention (AKA different coding of the query). Tuning SQL, or coding for optimum SQL efficiency, requires that the programmer understanding the logic of the Optimizer’s choices and to change her/his submitted code in a way that allows the Optimizer to make better choices. Basic to the tuning process is the understanding of what the Optimizer did when it ran a “too slow” SQL Query. _Method and _Tree show much of this information.

Unfortunately _Method and _Tree options produce large amounts of output. This paper will present an overview of the subject and the reader is encouraged to work thorough the annotated logs in the appendix of this paper (included on the on the CD). The annotated log files, in the appendix, is one of the larger, more detailed, collections of annotated _Method and _Tree output.

This is part of a planned series of papers on the SAS SQL Optimizer. The papers will build on each other and, hopefully, create a coherent body of knowledge on the Optimizer.

THE LIST OF KNOWN MESSAGES AND BASIC SQL PROCESSES

_Method and _Tree will show much about how a query executed, but use many abbreviations to indicate SAS SQL processes. In order to understand the cost penalties/implications of a particular execution plan, we must understand the abbreviations and some of the details of the processes used by SAS SQL.

A comment on, and apology for, the incomplete nature of this paper is in order. It is only too likely that the Optimizer currently has more subtlety than has been uncovered by the author and presented here. Additionally, SAS Inc. is constantly improving its products. As time goes by, the Optimizer will only become more powerful and more subtle and this paper will only become more incomplete.

TOP-TO-BOTTOM READ:

In certain situations SQL will perform a top to bottom read of the data. It will often do so in a query that does not have a where clause, or has a where clause without a usable index. The query might not have a usable index because 1) there is no index on the variable in the where clause, or 2) the code/syntax in the where clause might have prevented the Optimizer from using the index. SAS has put lots of time into making a top-to-bottom read fast and has been successful. Each read of a single observation is fast, but when millions of observations must be read, the total time (time=Number of Obs. * seconds per observation read) can still be unacceptably long.

Here are two examples of queries that will be executed in a top to bottom read.

Proc sql; Select * from dsn;	Proc sql;/*no index on sub*/ Select * from dsn WHERE SUB="001";
---------------------------------	---

EQUIJOIN

An equijoin is the name for a join that has an equality in the where clause. SAS has not implemented the code that an academic might consider a "true equijoin" however it has some fast techniques that the Optimizer can use on equality relationships. As a result, the SQL Optimizer is often able to process equijoins quickly. Some SUGI/NESUG articles have shown techniques to speed up queries by converting them to equijoins.

Proc sql;/*Equijoin*/ Select * from dsn WHERE SUB="001";	Proc sql;/*not an Equijoin*/ Select * from dsn WHERE SUB LE "001";
Proc SQL; /*Equijoin*/ SELECT L.SUBJID, L.NAME, R.AGE FROM LeftT as L, Right as R Where L.subjid=r.subjid;	Proc SQL; /*not Equijoin*/ SELECT L.NAME, L.Age, R.Name, R.Age FROM Left as L, RightT as R where L.Age GE R.Age;

SQL, whenever it can save time, delegates operations to lower SAS processes. In simple cases of an equijoin, like where Age=5, the where clause can be delegated to the data engine. The data engine will perform this equijoin (think of it as passing only obs where Age=5) and pass the result to SQL. If the where clause contains code like where age*12=60, all observations would be brought into SQL. The multiplication and filtering would happen in SQL. The data engine is usually not smart enough to perform multiplication/division and similar operations.

CARTESIAN PRODUCT OR STEP-LOOP JOIN

Consider merging two tables (LeftT and RightT) in a SQL from clause, one table on the left (LeftT) of the comma and one on the right (RightT). LeftT stands for the table on the left of the comma and RightT stands for the table on the right of the comma. Cartesian products and step loops are related merging processes and the SQL Optimizer employs them as a last resort. They are slow and very much to be avoided.

For these processes, a page of data is first read from the left table and then as much data as can fit in memory is read from the right table. All merges are made (between any observation in the page from LeftT and all observations in memory that came from RightT). The results are then output. Then the observations from RightT are flushed from memory and a new read of RightT pulls in as many observations as can fit in memory. All possible matches are made (between any observation in the page from LeftT and all observations in memory that came from RightT) and results are output. This process continues until SQL has looped through all of the table RightT.

Then SQL takes a step in the left table and reads a new page of data from LeftT into memory. The process of looping through right table is repeated for the second page from the left table. Then a third page is read from the left

table and the looping through the right hand table is performed again. The process continues until all the observations in the left table have been read and matched against every observation in the right hand table.

If there is a where clause in the query, it will be applied before the observations are output. SQL joins that are not based on an equality are candidates for the step loop process processes and are therefore to be avoided.

<pre>Proc sql; Select * from LeftT,RightT;</pre>	<pre>Proc sql;/*no index on sub*/ Select * from LeftT, RightT Where LeftT.sub < RightT.sub;</pre>
--	--

Both examples above are likely to be executed using a step-loop join (depending on system conditions and the decision of the SQL Optimizer).

INDEX JOIN

Even if an index exists on a variable in the where clause, the where clause (the code used for the query logic) can “prevent” the Optimizer from using the index. Additionally, even if the index exists and the code in the where clause does not disable the index, the Optimizer may decide not to use the index. The decision logic for the Optimizer is complex, but a firm index rule is: if index merge will return more than 15% of the indexed file to the result file, the index will not be used. In this case, there is a faster way to perform the query and the Optimizer will look for it. The Optimizer (or the data engine) can access metadata and will estimate output file sizes.

In the case of a one file select (see left box below) the Optimizer checks the percentage of observations that come from PA and their distribution in the file. If the Percentage is small, SQL reads the index file on the indexed variable “state”. It locates the desired level(s) of “state” in the index and reads, from that index-observation, the hard drive page number(s) that contain observations where state=“PA”. In order, the page number(s) are retrieved from the index, page number(s) are passed to the disc controller and data is sent back to the CPU. As each page of observations is received by the CPU, it is parsed and observations with state=“PA” are passed back to the working file for the query. SAS keeps track of the most recently read page, as a technique to minimize disk reads. If a page contains several observations that meet the where clause (e.g. state=“PA”), a new page will not be read from the hard drive until all the observations in that page from PA have been sent to the query working file.

In an index merge of two files (see right box below), an index exists and the where clause code must be written in a way that allows the SQL Optimizer to use it. In the appendix, several different queries are used to test how indexes are used by the SQL Optimizer. The basic process for index join is that one file is read from top to bottom and the matching observations in are read from the other file via an index lookup.

<pre>Proc sql;/*index on state*/ Select * from LeftT Where state=“PA”;</pre>	<pre>Proc sql;/*L_I_sub indexed in LeftT*/ Select * from LeftT , RightT Where Lft_tbl.L_I_sub = Rgt_tbl.subj</pre>
--	--

For more of an explanation of the process in the right box above; Assume two tables, LeftT and RightT, with an index on the variable L_I_sub in LeftT. The basic index merge process is that we are reading the right file from top to bottom and using an index to find observation(s) with matching subject Ids from left.

- 1) Select the next observation from RightT (at start, this is the first observation in the data set)
-if end of file, stop
- 2) Pass subj from RightT to the SAS index subroutine
- 3) Search the index on the variable L_I_sub in the file LeftT for the value of subj
- 4) If found, go to disk and return the required fields for that observation from LeftT
- output and go to 1)
- 5) If not found – go to 1)

If the information required by the query is in the index file itself the Optimizer will simply access the index, and not proceed to access the file associated with the index. This situation usually arises in queries when the where clause tests for existence of a match in another file (where a value in the variable subj in RightT is also in LeftT and Left_T has an index on subj). The fact that the Optimizer has automated this speed feature is just one indication of the level of detail that has gone into the designing and programming Optimizer, and of the difficulty of understanding it.

HASHING JOIN

Hashing can be a very fast technique and is automatically considered by the Optimizer. SQL hashing has been installed since V6.08 but since the Optimizer evaluates, and implements, the hashing technique without the programmer's intervention, its existence is not well known. General information on Hashing can be found in articles by Dr. Paul Dorfman in SUGI and NESUG online proceedings.

Hashing will not be considered as a join technique unless certain conditions are met. The SQL Optimizer accesses metadata on the file and "takes a good guess" at the size of the files it needs to join. After removing unneeded observations and variables, the Optimizer checks to see if 1% of the smaller of the two files being joined will fit into one memory buffer. If the smaller file appears to fit, the Optimizer will attempt a hash join. If the smaller file is too large, in relation to the buffer size, hashing will not be selected. A programmer can influence the Optimizer's choice of hashing as a merge technique by manually changing the buffer size with a SAS option.

SQL performs a hash join in the following way. The SQL Optimizer determines which of the two tables is smaller (after keeping only the appropriate variables from the select statement) and checks that smaller table size against the buffer. If the table meets the size criteria, it is loaded into a tree-like structure (a hash table) in memory. The structure of the hash object and the fact that is memory resident, allows for very fast searching. Then SQL processes the large file, from top to bottom, and for every observation that satisfies the where clause, it performs a HASH table lookup for the observation. Details of "data step" hashing are given in an article titled "An Annotated Guide: Resource use of common SAS Procedures" in the NESUG 2004 proceedings.

HASH & INDEX & WHERE USE

The Optimizer will dynamically adjust to new information. In certain situations it will switch from one join method to another-in the middle of execution - and create a hybrid join method. One such example is in the appendix (see examples 9A to 9D) and indicates that SQL simultaneously used a hash join and an index to produce the results of the query.

What happened is that, after tentatively trimming rows and columns from both files, the Optimizer estimated that 1%, of the smaller of the files being joined, would fit in a buffer. This is a strong hint/instruction for the Optimizer to use a hash join and so SQL loaded the smaller table into a hash table.

The Optimizer, as the hash table was being created, counted the number of unique key-variable values being loaded into the hash table. The number of unique values loaded into the hash table was found to be a small number (maybe below 1024) and the Optimizer dynamically changed the plan to take account of this information. In general, if there are fairly few unique values key in the hash table, the Optimizer will take the values from the hash table and use them to build an "in" phrase for a where clause (e.g. where state in("PA", "TX")).

SQL will then use the where clause to select observations but the Optimizer will again re-evaluate it's options in light of currently known information. The method selected for the join can be a top to bottom read or an indexed lookup. This adjustment of code to the details of a particular query is complex, dynamic and automatic. It can be seen in examples 9A to 9D.

SORT MERGE JOIN

This is similar, but not identical to, the data step merge. Under certain situations, the Optimizer determines that the fastest way to execute the query is to sort the tables and process both tables from top to bottom, using a merge that is similar (only similar) to that used by the DATA Step. This SQL merge will produce a Cartesian product, unlike the data step merge, and it does this by looping within the BY-group variables. SQL processes a page of data from the left table and loops through the appropriate by group right table.

GROUPING

Grouping, or aggregating observations, is a multi-step process and can take some time.

SELECT

Select statements specifies variables in the final data set. The optimizer, as part of it's run plan, creates "temporary working files" that are called result sets. Select logic is executed dynamically and as early as possible, keeping results sets small.

In the code below, only the variables name, age and sex are all brought into the original SQL query space (AKA result set). This initial removal of variables (height, weight) is handled by the data step engine and functions much like a keep option on a data set. Height and weight never become part of a SQL result set.

Observations with sex NE "M" are filtered out during the initial read of the data and that variable (sex) is eliminated from the query space, by the Optimizer, after completion of the read. After the completion of the first read, the result

set contains name and age. The result set is then sorted (by calling Proc Sort) by age. After the data is sorted, age is not required and the Optimizer eliminates that variable from the result set.

```
proc sql _method _tree;
create table lookat as select name from sashelp.class
where sex="M" order by age;
```

As the above explanation details, the Optimizer has automated “Good Programming Practices” and eliminates both variables, and observations, as soon as it can.

HAVING

The having statement is very useful to SQL programmers but requires that SQL perform several steps. Please examine the code below.

```
proc sql _method _tree;
title "this illustrates a having clause";
select name, sex, age
from sashelp.class
group by sex
having age=max(age);
quit;
```

In the code above, SQL must process the entire table sashelp.class, reading in the only the three variables in the select clause. SQL stores the observations in a result set (a temporary table that the SQL Optimizer directs be created). Then SQL makes a pass through the temporary table to find the max age, within each group, and tries to store that information in a “pipe line”. As a speed/storage technique, the Optimizer tries to avoid the creation of temporary files.

Sometimes applying a “having” requires additional passes through the working data set (AKA the result set) to check the having criteria against each observation. If there is no Note in the log mentioning re-merging, the Optimizer was able to produce the desired result in one pass using “pipe lines” to apply the having criteria. The query above produces the log below, where the word remerging indicates an additional pass was required to find the observations with the maximum age for each gender.

NOTE: The query requires **remerging** summary statistics back with the original data.

NOTE: SQL execution methods chosen are:

```
sqxslct
      sqxsumg
      sqxsort
      sqxsrc( SASHELP.CLASS )
```

DISTINCT

SQL usually implements distinct-ing by passing the result set to a Proc Sort with a nodup/nodupkey option. Distinct-ing is usually performed late in the query process as an additional pass through the data. This is a sorting and sorts are to be avoided because they take both time and space.

Under certain conditions (see examples 6A - 6D), the Optimizer can eliminate this last pass through if the query is distinct-ing a variable that has a unique index. The elimination of the distinct-ing saves time. There is an example of this, in compare-and-contrast format, in the appendix. It would be appropriate to use this as an example of how much effort has been put into making the Optimizer produce fast code. This situation does not happen often but the Optimizer has logic to help the programmer when it occurs.

UNCORRELATED SUB-QUERY

The code in an uncorrelated sub-query is processed just once by the SQL Optimizer. The results of this first evaluation are held in a result set until they are needed by the outer query. The code below is an example of an uncorrelated query. The result sets L and R are created once and accessed many times.

```
proc sql _method _tree;
  *title show inner join merge using a comma;
  create table ex5 as
  select coalesce (l.name, r.name) as name
    FROM ( select Name, Height from sashelp.class) AS L
        ( select Name, sex from sashelp.class) as R
  where l.name =r.name;
```

Outer

Uncorrelated
Inner Queries
or sub-queries

CORRELATED SUB-QUERY

A Correlated sub-query is processed differently from an un-correlated sub-query and again shows the power of the SQL Optimizer. A correlated sub-query uses information from each of the observations in the outer query to drive a “look up” process (a SQL query) against another table. In the worst case, SQL might have to execute the “look up process” for each row in the outer table. The look-up might have to be executed for every observation in the outer query but the Optimizer creates code that avoids that, whenever possible.

In the query below, as SQL processes each observation in the outer query, it seems to be passing the gender to the subquery and asking the subquery to find the maximum age for the current (in outer) value of gender. In fact, the Optimizer will process the sub-query for the first observation and store the results in a result set that is both temporary and indexed.

When additional observations from “outer” are processed the Optimizer first tries to find the needed information (in this case, has SQL calculated max age for that sex before) in the temporary, indexed result set. If it can find the required information in the temporary indexed result set, it takes information from the temporary indexed result set and does not execute the sub-query. If the information is not in the temporary indexed result set, SQL will run the sub-query, pass results to the outer query and then add the results of the query to the indexed result set. The temporary indexed result set grows in size as unique values of the equality variable are found in the outer file. When the query is done, the temporary indexed result set is deleted from the work library. The query below produces the _method output that follows it.

```
_method output that follows it.
proc sql _METHOD _TREE;
  TITLE "A SIMPLE CORRELATED QUERY";
  select * from sashelp.class as Outer
  where Outer.AGE =
    (select Max(age) from sashelp.class as inner
     where outer.sex=inner.sex);
quit;
```

Outer Query

Correlatedsub- query


```
NOTE: SQL execution methods chosen are:
      Sqxslct
      Sqxfil
      sqxsrc( SASHELP.CLASS(alias = OUTER) )
NOTE: SQL subquery execution methods chosen are:
      Sqxsubq
      Sqxsumn
      sqxsrc( SASHELP.CLASS(alias = INNER)
```

SIMPLE SUBSETTING LOGIC

For a one file SQL query, as shown below, the Optimizer must first decide if it should push the where down to the

data engine or handle it in SQL. The Optimizer's second decision is between using an index or a top-to-bottom read. The Optimizer has access to metadata on the file. The Optimizer considers both the percent of the file that will be returned and the distribution of the values in the file as it creates a plan to minimize run time.


```
Proc sql; /*index on state*/
Select * from LeftT
where state="PA";
```



The Optimizer has the ability to examine the metadata for the table and determine information useful for running the query. The metadata not only tells if there is an index on the variable in the where clause, but allows the Optimizer to determine both 1) what percent of the file will be returned by the where and 2) how the values are distributed through the file. This information lets the Optimizer make intelligent decisions on how to quickly access data.

MERGE LOGIC

The approximate logic for selecting a particular join is shown below (DSET means Data Set or table). Unfortunately this flowchart, like the one above, is a working model, rather than a definitive description of the Optimizer logic. The author's only consolation, is that whatever the current logic is, SAS Inc's commitment to improving it's product means that the current Optimizer will soon be replaced with one with more effective and subtle decision rules.



It is known that a where clause containing variables from two, or more, files can not be passed to the data engine.

OUTPUT FROM _METHOD

Method sends little output to the log. Below is typical output. It is important to note that an "indentation level" indicates the existence of a result set (working table, or temp file). Output in the appendix has been annotated. The query

```
proc sql _method;
  title Ex1 - show * select;
  create table ex1 as select * from sashelp.class;
```

produces the following _Method output in the log.

NOTE: SQL execution methods chosen are:

```
Sqxcrt
sqxsrc( SASHELP.CLASS)
```

A table of abbreviations is required to interpret the _Method output. Note that all abbreviations shown start with SQX. This prefix stands for “SQL Execution” code. Below, please find the abbreviations I have been able to collect while investigating _Method and a short explanation of the abbreviations. It is likely that more exist.

Name code	Description
SqxCRTA	Create table as select
SqxSLCT	Select
SqxJSL	Step loop join (Cartesian)
SqxJM	Merge Join
SqxINDX	Index Join
SqxHASH	Hash Join
SqxSORT	Sort
SqxSRC	Source rows from table
SqxFIL	Filter rows
SqxSUMG	Summary stats with group by
SqxSUMM	Summary stats with NO group by

OUTPUT FROM _TREE

The Optimizer creates a program, a multi-step program, and the output from tree can go on for pages. Understanding the run plan requires information from both _Method and _Tree. Below is manually annotated (the red numbers in parenthesis) _Tree output from the query above. Since _Tree output is quite complex, and explained in the annotated logs, only a cursory explanation will be given here.

Tree as planned.

```

      /-SYM-V-(class.Name:1 flag=0001)
      /-OBJ----| (5)
      | (3)    |--SYM-V-(class.Sex:2 flag=0001)
      |        |--SYM-V-(class.Age:3 flag=0001)
      |        |--SYM-V-(class.Height:4 flag=0001)
      |        \-SYM-V-(class.Weight:5 flag=0001)
      /-SRC----|
--SSEL---| (2)  \-TABL[SASHELP].class opt=''
(1)          (4)

```

Processing Sequence is:
 1) rightmost level first
 2) from bottom to top inside a level
 3) at the the top of each level, summarize what will be passed to the level to the left
 4) when a level is done, then step to the left one level

The query is processed, in levels, from right to left, and within “levels” from bottom to top. The idea of “passing data and information/instructions to the level to the left” is a useful mental concept. At the top of each level is a summary of the variables being passed to the level to the left. The result set being passed to the left is object (3) and it consists of variables (5).

In the output above,

- (4) SQL reads SASHELP.Class with no options processes by SQL
 (keep, drop etc are processed by the data set option processor).
- (5) SQL reads variables from a SAS dataset (in sashelp) called class. The variables are:
 Name: which is variable 1 in the data set class
 Sex: which is variable 2 in the data set class

 Weight: which is variable 5 in the data set class
- (3) is a summary of variables being passed to the level to the left.
- (2) indicates that the branches to the right describe a data source
- (1) indicates that this is a select type query (SAS developers can see other types)

Below, please find the abbreviations I have been able to collect while investigating _Tree and a short explanation of the abbreviations. Thanks to people at SAS for help with explanations.

	Abbreviation	This abbreviation can be found in Example Query Number found in the Appendix	Process
	ADIV	15	Divide
	AGGR	7, 8, 11	This indicates an aggregation, but SQL does several types of aggregation. This is associated with an aggregation operation like “select sex, sum(x) as totl group by sex”, or “select name, Min(x) as smallest” Sometimes processing the AGGR requires a separate pass through the data set (look for re-merging note in the log as an indicator of a separate pass) and some times it does not.
	AMUL	14	Multiply – A rithmetic M ultiplication
	ASC	3, 6A, 6C, 8, 13	Sort in ASC ending order
	ASGN	4, 5, 7	Assign. Create a new variable or Assign a value to a new variable. If the SQL code is “select sum(x) as totl” X will be summed and the result assigned to a variable named “totl”. If the programmer names the new variable, the name will be used in output and the variable will be easy to identify/track through the output. If the programmer does not name the variable, it will be numbered and can be tracked via the number. It is suggested that created variables be named as shown below Select coalesce(l.name, r.name) as Cname , r.age as Rgt_age
	CEQ	4, 5, 8, 13	This indicates a logical instruction to be passed to the level to the left, where it is executed. CEQ means “check if these 2 leaves are equal” CEQ is the symbol for both numeric and character equality testing.
	DESC	13	The sort, on the level to the left, should be in descending order. This DESC code is “information passed to the left” on the output. The descending sort is performed (files are created and time is spent) in the sorting in the level to the left of the DESC.
	Dlist		List of variables with distinct values that are participating in an aggregation/summation/grouping. See Slist and Tlist. Distincting gets rid of duplicates and Dlist reports on the distincting process. Variables on a dlist have to have duplicates removed before you can apply the aggregation.
	Empty	3, 4, 5, 6B, 6C, 7 , 8	This is a place holder in the output. The Optimizer has the capacity to do additional operations at this point-operations that were not performed.
	FCOA	4, 5,17,18,	This indicates a function, like a SAS function, of type

			coalesce.
	FIL	23	FILter is used to handle the situations that can not be handled by the data engine that is feeding into it. Fil indicates that SQL applied an additional predicate late in the processing. The clearest example of this is Ex 23. The where clause contains a variable that is not in the source data set (age_mo=age*12). It first must be created by SQL and then the filter predicate can be applied by SQL.
	Flags (class.Sex:2 flag=0001)		Flags are for developers and are also used internal to SQL. They are beyond the scope of this paper but, as one example, (001) means that the variable is used higher up in the SQL processing.
	JTAG	17, 17A, 19, 19A, 19B, 19C, 19D, 20A, 20B, 20C, 20D, 21A, 21B, 21C, 21D	This is a code that tells what type of join was applied. JDS=1 indicates a left join, JDS=2 Indicates a right join and JDS=3 indicates a full join
	FROM	4, 5	From indicates that data sources are being passed to higher (more leftward) process.
	GRP	8,	Group –is a multi-step process and can take some time to perform
	JOIN	4, 5	This indicates the SQL did a join, but not which type
	LAND	12	A Logical AND should be performed. This is usually in a where or a select.
	LITC	Not shown	This indicates the use of a Literal Constant (character string) as in: where state="PA"
	LITN	13,14 ,	This indicates the use of a Literal Number (ie numeric constant) as in: Where age LT 12
	OBJ	1,2, 3, 4, 5, 6A, 6B, 6C, 7, 8, 13, 14,	This indicates the existence of an object, or result set. Obj indicates a description of columns in a result set that is passed leftward for more processing. Objects come from data sets or lower objects.
	OBJE	4, 5, 7,14, 15,	This indicates the existence of an evaluated object, the result of an assignment of a value to a variable. An OBJE is a variable that is typically added/merged to another object (result set) at the same level. When a programmer codes Select sex, max(age) as Mage The value of maximum age will be assigned to Mage and mage will be, for a brief moment before the merge into the result set for that level, an OBJE. See example 4.
	ORDR	3, 6A, 7, 8,	Order By, On the tree, contains ordering information that is passed to the sort to the left. Order is information and not an instruction. It is not, and does not create, a result set. The sort, to the left of the ORDR, creates the result set.
	OTRJ	17, 18, 19	This stands for any of the following joins: Left join, right join, full join. OTRJs are paired with JTAGs and the JTAG indicates the type of join.
	SLST	7, 8,15	Indicates a list of things that are involved in an aggregation/summarization. This is a "Not distinct" List of things that "participate" in the summarization – See Dlist and Tlist
	SORT	3, 7, 8, 13	Sort, at this level, in the order described to the right.

	SRC	1,2, 4, 5, 6A, 6B, , 6C,7 , 8,	Information to the right of this is a data source. Typically an object is being passed to the left.
	SSEL	1,2, 3, 4, 5, 6A, 6B, 6C,7 , 8	SSEL indicates that the query is a Select query. Not all queries are of type select. There are modify queries and drop queries and others.
	SUBP	25	This indicates an input to the subquery. It is a parameter passed to the subquery
	SYM-A	4, 5, 7, 14,	This identifies a variable as an assigned/created variable- one that did not get read from a data set. A SYM-A is the result of Assigning a value to a variable through a calculation or function as in: Select name ,Age_yr=age*12 as YRS
	SYM-G	7, 8	A SYM-G is the result of assigning a value to a variable through a G rouping calculation or function as in: Select name ,sex, Max(age) as Mage Group by sex;
	SYM-V	25	This identifies a variable as a having been read from a data set.
	Sym-v lib.name Flag=0001	1, 2, 3, 4, 5, 6A, 6B, 6C, 7, 8,14,	This shows a combination of abbreviations as they might appear in the log. A variable was read from a source table (SAS data set lib.name). See Flag above
	Table [lib].fname opt=" "	1, 2, 3,4 , 5 , 6A, 6B, 6C, 7, 8	This shows a combination of abbreviations as they might appear in the log. Tables are identified with a two part name. See opt=
	TLST	7, 8,15	Indicates a summarization. A temp List of things that "participate" in the summarization. See Dlist and Slist
	UNIQUE	6A, , 6C,	This is the message to the log when select distinct is coded. Creating unique values is usually implemented by passing the current result set to Proc Sort with a nodup/nodupkey
	opt=" "		SQL allows data set options inside the Query. An example might be From class(keep=name age) Some of these options are processed by Base SAS and some are processed by SQL. If an option is processed by SQL, it will show up in the opt=" " note.

CONCLUSION

The SQL Optimizer is a tremendous help to programmers, allowing them to write very efficient queries with absolutely no thought. The amount of work that the SQL Optimizer does, independently of programmer input and totally behind the scenes, is amazing.

In some cases the performance may be improved by re-coding the query and passing the Optimizer different instructions. These issues will be explored in future papers.

If SQL performance is causing problems, knowing what the Optimizer created for a plan of execution is essential if the programmer want to attempt to improve performance. _method and _tree allow the programmer to see how SQL executed the query and to see the effects of her/his programming changes.

REFERENCES

TS553–SQL Joins –the Long and the Short of it, by Paul Kent – available on the SAS web site
TS320-Inside PROC SQL’s Query Optimizer, by Paul Kent – available on the SAS web site

Church (1999), Performance Enhancements to PROC SQL in Version 7 of the SAS® System
Performance Enhancements to PROC SQL in Version 7 of the SAS® System, Proceedings of the Twenty-fourth Annual SAS Users Group International Conference”, 24 , paper 51

Kent, Paul (1995) “SQL Joins – The Long and The Short of IT Proceedings of the Twentieth Annual SAS Users Group International Conference, Cary, NC: SAS Institute Inc., 1995, pp.206-215.

Kent, Paul (1996) “An SQL Tutorial – Some Random Tips”, Proceedings of the Twenty-First Annual SAS Users Group International Conference pp. 237-241.

For non-SAS explanations of SQL execution, see the article by Dan Hotka at www.odtug.com

ACKNOWLEDGMENTS

The author wishes to thank the Paul Dorfman, Sigurd Hermansen, Kirk Lafler, and Paul Sherman for their contribution to the SAS community on SQL and for inspiring this paper. Thanks to Paul Sherman for his review and comments.

Thanks for the help from SAS institute, especially help from Paul Kent and Lewis Church.

CONTACT INFORMATION

Your comments and questions are valued and encouraged. Contact the author at:

Russell Lavery Email: russ.lavery@verizon.net
or c/o Numeric, LLC
5 Christy Drive
Bldg.2 Suite 107
Chadds Ford, PA, 19317 (610)642-0700

SAS and all other SAS Institute Inc. product or service names are registered trademarks or trademarks of SAS Institute Inc. in the USA and other countries. ® indicates USA registration.

Other brand and product names are trademarks of their respective companies.

The Appendix Follows

****Example 2;**

```
proc sql _method _tree;
title Ex2 - show basic select;
create table ex2 as
  select Name, Height
  from sashelp.class;
```

Instead of select *, specify the variables

NOTE: SQL execution methods chosen are:

Sqxcrt (1) apply the selection criteria to observations???
sqxsrc(**SASHELP.CLASS**) (2) this indicates a source for the data

Tree as planned.

The position number of the variable in the data set

```

      /-SYM-V-(class.Name:1 flag=0001)
    /-OBJ----| (5)
  /-SRC----| (3) \-SYM-V-(class.Height:4 flag=0001)
--SSEL---| (2) \-TABL[SASHELP].class opt=''
(1)      (4)
```

Use this file (no options on the data set were processed by SQL)

Reading the above output shows:

(5) SYM-V shows us selecting only two variables from a SAS data set. (Note that the Optimizer is following good programming practice and only using variables it needs.) Then we see the two part variable name. The :1 means that name is the first variable in the data set (see contents above). The flag is a complex notation that is used by developers and internal processing. Except to say that the 1 means that the variable is required in higher level processing, flag is beyond the scope of this paper.

(4) This is an indication of the data source as SASHELP.class. This shown information to be passed to the left, where work is done. SQL accepts data set options(drop, keep, rename, etc) and can show in the opt= section. Some options are processed by base SAS and some by Proc SQL. If a data step option were processed by Proc SQL, it would be mentioned in the opt= ' '.

(3)OBJ, at the top of a level, summarizes variables being passed on to higher processing.

(2) SRC indicates a source of data, a result set that is being passed on to higher processing, or display. SRC indicates a "result set".


(1) SSEL indicates that this is a select query, not that variable selection happens at this point. A drop query, or modify table query would have a different string at this point.

****Example 3;**

```
proc sql _method _tree;
*title Ex3 - show basic select;
create table ex3 as
  select Name, Height
  from sashelp.class
  order by height ascending;
```

NOTE: SQL execution methods chosen are:

Sqxcrt (1) this indicates a selection of observations
sqxsort (2) this indicates a sorting of the data
sqxsrc(SASHELP.CLASS) (3) this indicates a source for the data



Reading the above output shows:

Note that ORDER is not a command to execute a sort. It is information about how the sort should be processed and is passed to the sort (2).

Order (11, 9, 5) is not an instruction to put observations in order. Order is information to be passed up to higher SQL processing. No processing happens in (11), (9) or (5). This section is simply telling SQL that the data will be sorted by ascending height- at a later time. The sort happens at (2).

(4) is a place holder. The optimizer could do amazing things here, but has not been requested to do so. :-)

(3) is a source of data. What is in that source is specified to the right.

(6A) is a summarization of the data being passed to the next higher level of processing. At this level the information about sort order (5) becomes important.

(2) sorting is done by the SAS Proc Sort. The order information is passed to SAS sort from (5).

(1) SSEL indicates that this is a select query, not that variable selection happens at this point. A drop query, or modify table query would have a different string at this point.

**Example 4;

```
proc sql _method _tree;
*title Ex4 - show inner join merge;
create table ex4 as
select coalesce (l.name, r.name) as name
  FROM ( select Name, Height from sashelp.class) AS L
      INNER JOIN
      (select Name, sex from sashelp.class) as R
    on l.name =r.name;
```


NOTE: SQL execution methods chosen are:

Sqxcrt (1) this indicates a selection of observations
 Sqxjsh (2) this indicates a hash join
 sqxsrc(SASHELP.CLASS) (8) this indicates a source for the data
 sqxsrc(SASHELP.CLASS) (9) this indicates a source for the data

Note: the Optimizer knows it does not need Height or Sex to produce the requested results. It will not bring them into the "QUERY SPACE". See (12) and (13).

We specify Inner Join not using comma. For illustration, extra variables are included.

Tree as planned.



(6) OBJE "documents" OBJect Evaluation logic.

It tells SQL that how to calculate/evaluate the variable.

(11) Coalesce two variables, put the result in an Assigned variable (type=SYM-A).

(10 , 5) this is not an process, like getting data (9). It is used to pass data to the SQL. When the files are merged, the values of names from the data sets (variables of type SYM-V) must be equal.

(9, 8, 4) take data from the sources to the right.

(3) This summarizes the data to be passed to the right.

What type of join?

(2) says that a join takes place. (5) says it is an equality join on name. Sqxjsh says that the join is a hash join.

Rule: The bottom file is the file that was loaded into the hash table.

NOTE that and OBJE, as well as an OBJ, can be passed up (7) started as (14). (14) was named, "name" (yeah, not too creative).


****Example 5;**

```
proc sql _method _tree;
*title Ex5 - show inner join merge using a comma;
create table ex5 as
select coalesce (l.name, r.name) as name
FROM ( select Name, Height from sashelp.class) AS L
,
( select Name, sex from sashelp.class) as R
where l.name =r.name;
```

NOTE: SQL execution methods chosen are:

Sqxcrt (1) this indicates a selection of observations
 Sqxjsh (2) this indicates a hash join
 sqxsrc(SASHELP.CLASS) (9) this indicates a source for the data
 sqxsrc(SASHELP.CLASS) (10) this indicates a source for the data

Tree as planned.



NOTE: Table WORK.EX5 created, with 19 rows and 1 columns.

(7, 12, 15, 16, 17) OBJE "documents" object Evaluation logic. It tells SQL how to calculate/evaluate the variable.

(16, 15) Coalesce two variables (both called name), put result in an Assigned variable (type=SYM-A)called name.

(11, 6) this is not an process, like getting data (9). It is used to pass data to SQL. When the files are merged in (2), the values of names from the data sets (variables of type SYM-V) must be equal.

(10, 9, 5) take data from the sources to the right.

(3) This summarizes the data to be passed to the right.

What type of join?

(2) says that a join takes place. (6) says it is an equality join on name. Sqxjsh says that the join is a hash join. The bottom file is loaded into the hash table.

NOTE that and OBJE, as well as an OBJ, can be passed up (8) started as (15). (15) was named, "name" (yeah, not too creative).

**** THIS IS AN EXAMPLE OF HOW MUCH THE OPTIMIZER DOES FOR US**

If we have a unique index on that variable that can satisfy the select, it eliminates the sort -> 6B

```
proc sql _method _tree;
```

```
create table ex6A as select distinct name, sex , age FROM sashelp.class;
```

Sqxcrt (1) this indicates a selection of observations

`sqxsrc( SASHELP.CLASS )` (4) this indicates a source for the data

[illegible]

```
data class_W_indx(index=(name/ unique));
set sashelp.class; run;
proc sql _method _tree;
    *title Ex6B - show distinct on unique indexed variable;
    create table ex6B as select distinct name FROM class_W_indx;
quit;
```

NOTE: SQL execution methods chosen are:

Sqxcrt (1) this indicates a selection of observations

Tree as planned.

```

--SYMB-V- (class_W_indx.Name:1 flag=0001)
--SRC----| /-OBJ----| (5)
--SSEL---| (2) \-TABL[WORK].class_W_indx opt=''
          (1) (4)

```

This Query did not even access the raw data.
It was able to get all the information it
needed from the index.

there is no need for string-parsing in the query.

****Example 6C *1 variable index, distinct on 2 vars. **;**

```
proc sql _method _tree;
  *title Ex6C -distinct on unique indexed variable when selecting multiple variables;
  create table ex6C as
    select distinct name ,sex
    FROM class_W_indx; quit;
```

We have a distinct index on ONE of the two variables in the select. The metadata can not be used to help us.


NOTE: SQL execution methods chosen are:

Sqxcrt (1) this indicates a selection of observations

Sqxuniq (2) this indicates a source for the data

sqxsrc(WORK.CLASS_W_INDX) (4) this indicates a source for the data

Tree as planned.



The metadata information (index information) does not contain enough information to allow the Optimizer to help. Create a new dataset with a new index that is "on" both the variables and then use it in Proc SQL.

****Example 6D ***Compound index and distinct *****;**

```
data class_W_Cmpindx(index=(nm_sex=(name sex)/ unique));
set sashelp.class;
run;
```

Create a unique compound index on Name and sex - In real life we'd worry about names like Pat that can be male or female (Patrick and Patricia), but not in this small data set.


```
proc sql _method _tree;
  *title Ex6D - show distinct on two vars with a unique compound index;
  create table ex6d as select distinct name, sex FROM class_W_Cmpindx;
```

NOTE: SQL execution methods chosen are:

Sqxcrt (1) this indicates a selection of observations

sqxsrc(WORK.CLASS_W_CMPINDX) (2) this indicates a source for the data

Tree as planned.




SQL reads the header information in the file and realizes that these obs. are all unique. It just prints the observations using the index itself as the source of the data. No mention of index use in the log, in _Method or in _Tree, but the unique index was sensed by the optimizer and used as the data source.

```
proc sql _method _tree;  *title Ex7 - Grouping and calculations;
    create table ex7 as    select sex , max(age) AS maxage, AVG(HEIGHT) AS AVG_HT
    FROM class_W_indv     GROUP BY SEX ;
```

Sqxcrt (1) this indicates a selection of observations
Sqxsumg (3) this indicates an aggregation - maps to AGGR in tree
Sqxsort (5) this indicates a sorting of observations
sqxsrc (**WORK.CLASS_W_INDX**) (4) this indicates a source for the data

Tree as planned.



In (26) SymG indicates that a grouping is involved. Stat=5 is a code indicating that an average is calculated for the groups. The left 0 in the (0,0) after age indicates the source of the data (position 0 in the slist). The right 0 indicates non-unique. There are three result sets here. (15) (5) and (3) are physical files.

****Example 9 exploring index use in a join *****;**

Create data for join. Note that in the data set Left, we have an indexed variable called **LIndxName. Note that in the data set Right, we have an indexed variable called **RIndxName**.**

When we merge the two datasets, left will always be to the left of the join indicator and right will always be to the right.

E.G. (From left as L, right as R)

*left and right both have an un-indexed variable called name and a variable with a two part name "L" or "R" and the suffix "IndName";

This naming convention will make it easier to understand if a variable has an index or not.

```
data left_class( drop= RIndxName index=(LIndxName))
  Right_class( drop= LIndxName index=(RIndxName));

  length name $ 13;
  set sashelp.class;

  do i=1 to 1200; /*expand file so we do not hash*/
    name=name||put(i,5.0);
    RIndxName=name;
    LIndxName=name;
    Output;
  end;

run;
```

The data step above just expands the class data set (it uses a loop to increase the file size in an effort to have files so large that 1% of the file is larger than the buffersize) to reduce the chance of using a hash in the merging. The concatenation is done so that there are unique values of name so that we can make a unique index.

```
create table hope as
select coalesce(l.name, r.name), l.sex, r.age
  From left_class as l inner join right_class as r
    on l.name=r.name; /*these are not indexed*/
```

NOTE: SQL execution methods chosen are:

- Sqxcrt (1) this indicates a selection of observations
- Sqxjm (2) this indicates a Merge join
- Sqxsrt (10) SORT
 - sqxsrc(WORK.LEFT_CLASS(alias=L))(14)observations
- sqxsrt (11) SORT
 - sqxsrc(WORK.RIGHT_CLASS(alias=R))(19)observations

Tree as planned.

[illegible]

Page 23 of 57

Page 25 of 57

NOTE: SQL execution methods chosen are:

Sqxjsh (2) this indicates a HASH join

HASH JOIN indicated “incorrectly”

```

anned.
/-OBJ----| (9)
(3)      \--SYM-V-(l.Sex:2 flag=0001)
          \-SYM-V-(r.Age:3 flag=0001)

/-JOIN---| (2)
          /-SYM-V-(r.RIndxName:7 flag=0001)
          (20)
          /-OBJ----| (14)
          (14)      \--SYM-V-(r.name:1 flag=0001)
                    \-SYM-V-(r.Age:3 flag=0001)

          /-SRC----| (10)
          (10)      \-TABL[WORK].right_class opt=''
                    (15)

          /-OBJ----| (21)
          (16)      \--SYM-V-(l.LIndxName:7 flag=0001)
                    \-SYM-V-(l.name:1 flag=0001)
                    \-SYM-V-(l.Sex:2 flag=0001)

          \-SRC----| (11)
          (11)      \-TABL[WORK].left_class opt=''
                    (17)

          --empty- (5)
          --CEQ----| (12)
          (6)      /-SYM-V-(r.RIndxName:7)
                    \-SYM-V-(l.LIndxName:7)

          --empty- (7)
          --empty- (7)

          /-SYM-A-(#TEMA001:1 flag=0031)
          /-ASGN---| (18)
          (13)      \-FCOA---| /-SYM-V-(l.name:1)
                    (22)
                    (19)      \-SYM-V-(r.name:1)

          \-OBJE---| (8)

```

HASH
 JOIN
 above
 and
 index
 join
 below

The indexed variables are
 passed up to the CEQ and
 dropped ASAP

Name is
Passed up,
& dropped
after
coalesce

The indexed variables are passed up to the CEQ and dropped ASAP

If there fairly few unique values key in the hash table, the Optimizer will take the values from the hash table and use them to build an "in" phrase for a where clause (e.g. where state in("PA", "TX")). The Optimizer dynamically adjusted the plan to use an index lookup to effect the merge.

****example 10 **when does select happen???*****;**

```
proc SQL _method _tree; /*EX10A the timing of the selects: variables is early */
select name , sex, age
  from sashelp.class
 order by sex, age;
```

NOTE: SQL execution methods chosen are:

Sqxsclt (1) this indicates a selection of observations

Sqxsrt (2) this indicates a Sort

sqxsrt(SASHELP.CLASS) (3) indicates a selection of observations

Tree as planned.

```

      /-SYM-V-(class.Name:1 flag=0001)
      | (7)
      | /-OBJ----|
      | (3) |--SYM-V-(class.Sex:2 flag=0001)
      |      \-SYM-V-(class.Age:3 flag=0001)
      |
      | /-SORT---|
      | (2) |
      |      |
      |      | /-SYM-V-(class.Name:1 flag=0001)
      |      | /-OBJ----|
      |      | (8) |--SYM-V-(class.Sex:2 flag=0001)
      |      |      \-SYM-V-(class.Age:3 flag=0001)
      |      |
      |      | /-SRC----|
      |      | (4) \-TABL[SASHELP].class opt=''
      |      |
      |      | /-empty--|
      |      | (9)
      |      | /-SYM-V-(class.Sex:2)
      |      | /-ASC----|
      |      | (13)
      |      | \-ORDR---|
      |      | (10)
      |      | (6) \-SYM-V-(class.Age:3)
      |      |      \-ASC----|
      |      |      (14)
      |      |      (11)
      |
      | /-SSEL---|
      | (1)

```

The Optimizer implements good programming practices. Variables are Selected Early. Un-needed variables (height, weight) are not brought into the SQL space.

(6, 10,11,13 &14) show information that is passed onto the sort (2). Sorting is done by Proc Sort.

NOTE: PROCEDURE SQL used (Total process time): real time 1.66 seconds cpu time 0.45 seconds

Note how the Optimizer only brings in variables it needs for the query.

***DYNAMICALLY Trimming Extra (not needed) variables As they become redundant ;**

```
proc sql _method _tree;title "EX10B timing of selects:of observations";
```

```
select name from sashelp.class
```

```
  where sex="M" order by age;
```

NOTE: The query as specified involves ordering by an item that doesn't appear in its SELECT clause.

NOTE: SQL execution methods chosen are:

sqxsclt

sqxsrt

sqxsrt(SASHELP.CLASS)

Tree as planned.

```

      /-SYM-V-(class.Name:1 flag=0001)
      | (7)
      | /-OBJ----|
      | (3) |
      |      | /-SYM-V-(class.Name:1 flag=0001)
      |      | /-OBJ----|
      |      | (6) \-SYM-V-(class.Age:3 flag=0001)
      |      |
      |      | /-SRC----|
      |      | (4) \-TABL[SASHELP].class opt=''
      |      |
      |      | /-empty--|
      |      | (9)
      |      | /-SYM-V-(class.Sex:2)
      |      | /-CEQ----|
      |      | (13)
      |      | (10) \-LITC('M')
      |      |
      |      | /-empty--|
      |      | (5)
      |      | /-SYM-V-(class.Age:3)
      |      | /-ASC----|
      |      | (14)
      |      |
      |      | /-ORDR---|
      |      | (6)
      |
      | /-SSEL---|
      | (1)

```

Note is from SQL

Age Not Needed in output and is not stored here

Age Needed for ordering

OBS. & Vars. are Selected Early. Base SAS Compares the value in sex to the literal character value "M"

Age Not Needed after ordering

NOTE: PROCEDURE SQL used (Total process time): real time 1.43 sec. cpu time 0.01 sec.

Note how the optimizer trims variables that it no longer needs. It needs to select for sex="M", and then drops sex. It orders by age, and then drops it as well.

*****example 11***The Having Clause *****;

```
proc sql _Method _tree;
title "EX11 this illustrates a having clause";
select name, sex, age from sasHELP.class
group by sex having age=max(age);
```

Remerging is associated with the AGGR and takes another pass through the data

NOTE: The query requires **remerging** summary statistics back with the original data.

NOTE: SQL execution methods chosen are:


Sqxslect (1) this indicates a selection of observations

Sqxsumg (2) Aggregate is associated with a having- requires a pass through the data

sqxsort (4) this indicates a Sort

sqxsorc(SASHELP.CLASS) (12) this indicates a selection of observations

Tree as planned.



The data was re-merged and a second pass was required to get the results.

*****example 12 **illustrates the and clause *****;

```
proc sql _Method _tree;
title "EX12 this illustrates a AND clause";
select name, sex, age from sasHELP.class
group by sex
having age=max(age) and sex="F";
```


Remerging is associated with the AGGR and takes another pass through the data

NOTE: The query requires **remerging** summary statistics back with original data.

NOTE: SQL execution methods chosen are:

Sqxslect (1) this indicates a selection of observations
 Sqxsumg (2) Aggregate is associated with a having- requires a pass through the data
 Sqxsort (4) this indicates a SORT
 sqxsrc(SASHELP.CLASS) (12) this indicates a selection of observations

Tree as planned.



Having is applied in the AGGR (2) and requires a pass through the sorted data and **remerging**.

The key to the having is the LAND (logical And) is at (5). We do not de-dupe in this query. Everyone having an age= max(age) gets passed on.

****example 13 **illustrates variable=literal & sorting**;**

```
proc sql _Method _tree;
title "EX13 this illustrates a = literal and sorting";
select name, sex, age from sashelp.class
  where age = 12 order by name desc , height asc ;
```

NOTE query as specified involves ordering by an item that doesn't appear in its SELECT clause.

NOTE: SQL execution methods chosen are:


Sqxslect (1) this indicates a selection of observations

Sqxsrt (2) this indicates a SORT

sqxsrc(SASHELP.CLASS) (4) this indicates a selection of observations

We sort by a var. (height) not in the final output

Tree as planned.



The checking of single observations against the criteria age=12 is done early – by the data engine. The optimizer wants to make tables as small as possible and will filter out observations (and variables) as soon as possible. Since obs with age NE 12 have been removed, SRC (4) is a small data set.

To increase speed, the Optimizer has eliminated both variables and observations.

```
proc sql _Method _tree;
title "EX14 this illustrates a = calculation";
  select name, sex, age*12 as agemo
    from sashelp.class ;
```

Sqxslct (1) this indicates a selection of observations
Sqxfil (2) this indicates the application of a predicate "late" in the process
sqxsrc(**SASHELP.CLASS**) (12) this indicates a selection of observations

```

ree as planned.

/-OBJ---|
| (7)
| (3) /-SYM-V-(class.Sex:2 flag=0001)
| \-SYM-A-(agemo:1 flag=0031)

/-FIL---|
| (2)

Fil
is the late
application
of a
predicate

--SRC---|
| (4)
--empty-|
--empty-|
--empty-|
--empty-|
--empty-|
| (5)

/-OBJ---|
| (8)
| (11) /-SYM-V-(class.Name:1 flag=0001)
| --SYM-V-(class.Sex:2 flag=0001)
| \-SYM-V-(class.Age:3 flag=0001)

\ -TABL[SASHELP].class opt=''
| (9)

/-ASGN---|
| (10)
| (12) /-SYM-A-(agemo:1 flag=0031)
| \-AMUL---|
| (13) | (14) /-SYM-V-(class.Age:3)
| \-LITN(12)

--SSEL---|
| (1)

\ -OBJE---|
| (6)

```

The optimizer passed Agemo to (2) but not age

Fil is the late application of a predicate

Since age is not in the output of the query, the data engine would “like” to not bring age into the result set. However, the data engine can not do the multiplication required for agemo. The Optimizer directs that the data engine bring in age so that SQL can use it in the multiplication. SQL calculates agemo and passes up the result. At the next higher level, the filter (2) on the variable age can be applied to reduce the size of the data set.

Page 31 of 57

*****example 15 *****This illustrates a division*****;

```
proc sql _Method _tree;
title "EX15 this illustrates a = division";
```

```
select name, sex, height/12 as Ht_feet
from sashelp.class
where age = 12
order by name desc ;
```

NOTE: SQL execution methods chosen are:


Sqxslet (1) this indicates a selection of observations

Sqxsrt (2) this indicates a SORT

Sqxfil (4) this indicates the application of a predicate late in the process

sqxsrc(SASHELP.CLASS) (9) this indicates a selection of observations

Tree as planned.



Note the ordering (12) and the division (ADIV).

Fil means that there is a filter that is applied here, that can not be applied earlier (to the right). The data engine handles the `age=12` restriction.

Since height is not in the output of the query, the data engine would "like" to not bring height into the result set. However, the data engine can not do division for `Ht_feet`. The Optimizer directs that the data engine bring in height so that SQL can use it in the division. SQL calculates `Ht_feet` and passes up the result. At the next higher level, the filter on the variable height, can be applied to reduce the size of the data set.

*****example 16 ***SUMMARY WITHOUT GROUPING*****;

```
proc sql _Method _tree;
title "EX16 this illustrates a = summary without grouping";
select name, sex, height, min(height) as shortest, max(height) as tallest
from sashelp.class order by name ;
```

NOTE: The query requires remerging summary statistics back with original data.

NOTE: SQL execution methods chosen are:


Sqxslect (1) this indicates a selection of observations

Sqxsrt (4) this indicates a SORT

Sqxsumm (6) this indicates summation without grouping—a summary of the whole table

sqxsrc(SASHELP.CLASS) (11) this indicates a selection of observations

Tree as planned.



Note the creation of the variables shortest and tallest in ASGNs (21) and (22) .)

```

*****
*****
*****
***** JOINS *****

```

```

data left_class( drop= RIdxName index=(LIdxName))
  Right_class( drop= LIdxName index=(RIdxName));
  length name $ 13;
  set sashelp.class;

  do i=1 to 1200; /*expand file so we do not hash*/
    name=name||put(i,5.0);
    RIdxName=name;
    LIdxName=name;
    Output;
  end;
run;

```

NOTE: There were 19 observations read from the data set SASHELP.CLASS.

NOTE: The data set WORK.LEFT_CLASS has 22800 observations and 7 variables

NOTE: Simple index LIdxName has been defined.

NOTE: The data set WORK.RIGHT_CLASS has 22800 observations and 7 variables.

NOTE: Simple index RIdxName has been defined.

NOTE: DATA statement used (Total process time):

real time	4.05 seconds
cpu time	0.27 seconds

Create two tables that allow us to examine joins done in several ways. We will examine joins with indexes in Left position, in right position, and in both positions.

No index. Use
Sort merge join

Coalesced Var.

Coalesced Var. –
note: the query did
not specify a name
for it

```
Proc SQL _method _tree;      title "EX17A Illustrating a right join" ;
create table hope as      select coalesce(l.name, r.name), l.sex, r.age
  From left_class as l right join right_class as r on l.name=r.name;
```

Sqxcrt (1) this indicates a selection of observations

Sqxsrt (11) this indicates a SORT

sqxsort (12) this indicates a SORT

Tree as planned.

No index.
Use Sort merge join

Coalesced Var.

JTAG
if ends in
1= left join,
2=right join
3=full join
and we
have
variations
on above

JTAG if ends in
1= left join, 2=right join 3=full join

Coalesced Var. –
note: the query did
not specify a name
for it

```

Proc SQL _method _tree;
title "EX17B Illustrating an INNER JOIN WITH COMMA with index on left table" ;
create table hope as
select coalesce(l.name, r.name), l.sex, r.age
  From left_class as l , right_class as r
 where l.LindxName = r.name;

```


NOTE: SQL execution methods chosen are:

```


Sqxcrt (1) this indicates a selection of observations
Sqxjsh (2) is a hash join
sqxsrc( WORK.LEFT_CLASS(alias = L) ) (10) indicates a selection of observations
sqxsrc( WORK.RIGHT_CLASS(alias = R) ) (11) indicates a selection of observations

```

Tree as planned.



L.name and R.name are coalesced (8, 13,18,19,22) and stored in a SYM-A variable (18) and kept as part of the output (9). While an index exists on the variable on the left side of the join, a hash join was selected by the optimizer.




```
Proc SQL _method _tree; title "EX18D INNER JOIN W/INNER JOIN Phrase w/index on BOTH tables";
create table hope as
select coalesce(l.name, r.name), l.sex, r.age
  From left_class as l inner join right_class as r
    on l.LindxName = r.RindxName;
```

NOTE: SQL execution methods chosen are:


Sqxcrt (1) this indicates a selection of observations

Sqxjhsh (2) is a hash join

sqxsrc(WORK.LEFT_CLASS(alias = L)) (10) indicates a selection of observations

sqxsrc(WORK.RIGHT_CLASS(alias = R)) (11) indicates a selection of observations

Tree as planned.



```
Proc SQL _method _tree; title "EX19A Illustrating a left join with no indexes";
create table hope as select coalesce(l.name, r.name), l.sex, r.age
  From left_class as l left join right_class as r on l.name = r.name; Use EET 10
```

Use LEFT JOIN Phrase!
No index to use

```
/-SYM-A-(#TEMA001:1 flag=0035)
```

```

/-OBJ----| (10)
| (3)      |--SYM-V-(l.Sex:2 flag=0001)
|          \-SYM-V-(r.Age:3 flag=0001)

```

/-0TRJ- - -
| (2)

```

                                /-SYM-V-(r.name:1 flag=0001)
/-OBJ----| (25)
| (15)    \-SYM-V-(r.Age:3 flag=0001)

```

```
| (11) | /-SYM-V-(r.name:1 flag=0001) |
|      | /-OBJ----| (38)
|      | | (26) \-SYM-V-(r.Age:3 flag=0001) |
```

```
|      |--SRC----|
|      |  (16)    \-TABL[WORK].right_class opt='
|      |--empty-  (27)
```

```
| (17) /-SYM-V-(r.name:1)|
| /-ASC----| (39)
| \-ORDR---| (28)
| (12)
```

```
|
|                                     /-SYM-V-(l.name:1 flag=0001)
|                                     /-OBJ----| (29)
| (19) \-SYM-V-(l Sex:2 flag=0001)
```

```

\ -SORT--- |
(12) | /-SYM-V-(l.name:1 flag=00)
| /-OBJ---- (40)

```

```
|          | (30)  \-SYM-V-(l.Sex:2 flag=0001  
|--SRC---|  
| (20)    \-TABL[WORK].left_class opt=''
```

```
| --empty-      (31)
| (21)          /-SYM-V-(l.name:1)
|              /-ASC----| (41)
```

```

(22)
/-SYN-V-(r.name:1)
(12)

```

equality check info passed up

JTAG: if ends
in
1= left join,
2=right join
3=full join

```
(6) \-SYM-V-(l.name:1)
|--JTAG(jds=1, tagfrom=1, flags=0)
```

```
| (7) --empty- /-SYM-A-(#TEMA001:1 flag=0031)
| (8) /-ASGN-- (23) /-SYM-V-(l.name:1)
```

	(14)	\-FCOA---	(33)
\-OBJE--	(24)	\-SYM-V-(r.name:1)	
(9)			

Page 44 of 57


```
Proc SQL _method _tree; title "EX19D Illustrating a left join with index on BOTH tables";
create table hope as select coalesce(l.name, r.name), l.sex, r.age
From left_class as l left join right_class as r on l.LIndxName = r.RIndxName;
```

Sqxcrt (1) this indicates a selection of observations
 Sqxjm (2) this indicates a sort-merge type of join
 Sqsxsort (11) this indicates a SORT
 sqxsorc(WORK.LEFT_CLASS(alias=L)) (16) indicates a selection of observations
 sqxsorc (12) this indicates a SORT
 sqxsorc(WORK.RIGHT_CLASS(alias=R)) (20) indicates a selection of observations

Use LEFT JOIN
 Phrase! Two vars
 have index

```
Tree as planned. /-SYM-A- (#TEMA001:1 flag=0035)
/---OBJ---| (10)
(3) |--SYM-V-(l.Sex:2 flag=0001)
\--SYM-V-(r.Age:3 flag=0001)
/---OTRJ---| (2)
(2) |
| /---SYM-V-(r.RIndxName:7 flag=0001)
| /---OBJ---| (25)
| (15) |--SYM-V-(r.name:1 flag=0001)
| \--SYM-V-(r.Age:3 flag=0001)
| /---SORT---| (11)
| | /---SYM-V-(r.RIndxName:7 flag=0001)
| | /---OBJ---| (34)
| | (26) |--SYM-V-(r.name:1 flag=0001)
| | \--SYM-V-(r.Age:3 flag=0001)
| |--SRC---|
| | (16) \--TABL[WORK].right_class opt=''
| |--empty-| (27)
| | (17) /---SYM-V-(r.RIndxName:7)
| | /---ASC---| (35)
| \---ORDR---| (28)
| /---SYM-V-(l.LIndxName:7 flag=0001)
| /---OBJ---| (29)
| (19) |--SYM-V-(l.name:1 flag=0001)
| \--SYM-V-(l.Sex:2 flag=0001)
| \---SORT---| (12)
| | /---SYM-V-(l.LIndxName:7 flag=0001)
| | /---OBJ---| (36)
| | (30) |--SYM-V-(l.name:1 flag=0001)
| | \--SYM-V-(l.Sex:2 flag=0001)
| |--SRC---|
| | (20) \--TABL[WORK].left_class opt=''
| |--empty-| (31)
| | (21) /---SYM-V-(l.LIndxName:7)
| | /---ASC---| (37)
| \---ORDR---| (32)
| (22)
|--empty-| (5)
| /---SYM-V-(r.RIndxName:7)
|--CEQ---| (13)
| (6) \--SYM-V-(l.LIndxName:7)
|--JTAG(jds=1, tagfrom=1, flags=0)
| (7)
|--empty-| (8)
| /---ASGN---| (14)
| (23) /---SYM-V-(l.name:1)
| \---FCOA---| (33)
| (24) \--SYM-V-(r.name:1)
| \---OBJE---| (9)
--SSEL---| (1)
```

JTAG
 if ends in
 1= left join,
 2=right join
 3=full join

Merge join has
 notation of
 JOIN
 Or OTRJ

 Merg joins can
 be used for
 inner join or
 outer joins

JTAG: if ends
 in
 1= left join,
 2=right join
 3=full join

R.Name and
 r.RIndxName
 are passed up,
 used & dropped
 after coalesce.

L.Name and
 Lindxname are
 passed up, used
 & dropped after
 coalesce & CEQ

equality check info passed up

Coalesced Var.

Use RIGHT JOIN Phrase!

L.Name and
L.LindxName are
passed up, used &
dropped after
coalesce & CEQ

r.name is
passed up, used
& dropped after
coalesce & CEQ

Merge join has notation of JOIN Or OTRJ

Merg joins can be used for inner join or outer joins

JTAG
if ends in
1= left join,
2=right join
3=full join

Merg joins
can be used
for inner
join or outer
joins

JTAG: if ends in
1= left join,
2=right join
3=full join

equality check info passed up

Coalesced Var.

(1) | | |
--SSEL--- | \-(9)OBJE-



Use RIGHT JOIN Phrase!

L.Name is passed up, used & dropped after coalesce & CEQ

r.name and R.RindxName are passed up, used & dropped after coalesce. & CEQ



JTAG
if ends in
1= left join,
2=right join
3=full join

Merge join has notation of JOIN
Or OTRJ

Merg joins can be used for inner join or outer joins



JTAG: if ends in
1= left join,
2=right join
3=full join

equality check info passed up



Coalesced Var.

Use RIGHT JOIN Phrase!

L.Name is
passed up, used
& dropped after
coalesce & CEQ

JTAG
ends in
= left join,
=right join
=full join

Merge join
has notation
of
JOIN
Or OTRJ

Merg joins
can be used
for inner join
or outer joins

JTAG: if end
in
1= left join,
2=right join
3=full join

equality check info passed up

Coalesced Var.

Page 51 of 57

*****example 21 ***** FULL JOIN *****;

```
Proc SQL _method _tree; title "EX21A Illustrating a full join with no indexes" ;
create table hope as select coalesce(l.name, r.name), l.sex, r.age
From left_class as l full join right_class as r on l.name = r.name;
```

Sqxcrt (1) this indicates a selection of observations

Sqxjm (2) this indicates a sort-merge type of join


Sqxsrt (11) this indicates a SORT

sqxsrc(WORK.LEFT_CLASS(alias=L)) (16) indicates a selection of observations

sqxsrt (12) this indicates a SORT

sqxsrc(WORK.RIGHT_CLASS(alias=R)) (20) indicates a selection of observations

Tree as planned. /-SYM-A- (#TEMA001:1 flag=0035)





Use FULL JOIN Phrase!

[illegible]


```

Proc SQL _method _tree; title "EX21D Illustrating a full join w/ index on BOTH tables" ;
create table hope as select coalesce(l.name, r.name), l.sex, r.age
From left_class as l full join right_class as r on l.LIdxName = r.RIdxName;
Sqxcrta (1) this indicates a selection of observations
Sqxjm (2) this indicates a sort-merge type of join
Sxsort (11) this indicates a SORT
sqxsrc(WORK.LEFT_CLASS(alias=L)) (16) indicates a selection of observations
sqxsort (12) this indicates a SORT
sqxsrc(WORK.RIGHT_CLASS(alias=R)) (20) indicates a selection of observations

```

Use FULL
JOIN Phrase!

```

Tree as planned.
/SYM-A-(#TEMA001:1 flag=0035)
/-OBJ----| (10)
| (3) |--SYM-V-(l.Sex:2 flag=0001)
| | \-SYM-V-(r.Age:3 flag=0001)
/-OTRJ---| (2)
| | /-SYM-V-(l.LIdxName:7 flag=0001)
| | /-OBJ----| (24)
| | | (15) |--SYM-V-(l.name:1 flag=0001)
| | | | \-SYM-V-(l.Sex:2 flag=0001)
| | /-SORT---| (11)
| | | | /-SYM-V-(l.LIdxName:7 flag=0001)
| | | | /-OBJ----| (32)
| | | | | (25) |--SYM-V-(l.name:1 flag=0001)
| | | | | | \-SYM-V-(l.Sex:2 flag=0001)
| | | |--SRC---| (16)
| | | | \-TABL[WORK].left_class opt=''
| | | |--empty-| (26)
| | | | (17) /-SYM-V-(l.LIdxName:7)
| | | | /-ASC----| (32)
| | | | \-ORDR---| (27)
| | | (18) /-SYM-V-(r.RIdxName:7 flag=0001)
| | | /-OBJ----| (28)
| | | | (19) |--SYM-V-(r.name:1 flag=0001)
| | | | | \-SYM-V-(r.Age:3 flag=0001)
| | | \-SORT---| (12)
| | | | /-SYM-V-(r.RIdxName:7 flag=0001)
| | | | /-OBJ----| (33)
| | | | | (29) |--SYM-V-(r.name:1 flag=0001)
| | | | | | \-SYM-V-(r.Age:3 flag=0001)
| | | |--SRC---| (20)
| | | | \-TABL[WORK].right_class opt=''
| | | |--empty-| (30)
| | | | (21) /-SYM-V-(r.RIdxName:7)
| | | | /-ASC----| (33)
| | | | \-ORDR---|
| | | (22)
| | | (5) /-SYM-V-(l.LIdxName:7)
| | | /-CEQ----| (13)
| | | | (6) \-SYM-V-(r.RIdxName:7)
| | | | \-JTAG(jds=3, tagfrom=3, flags=0)
| | | | (7)
| | | |--empty-| (8)
| | | | /-ASGN---| (14)
| | | | | (22) /-SYM-A-(#TEMA001:1 flag=0031)
| | | | | /-SYM-V-(l.name:1)
| | | | | \-FCOA---| (31)
| | | | | (23) \-SYM-V-(r.name:1)
| | | \-OBJE---| (9)
| | (1)
--SSEL---|

```

JTAG
if ends in
1= left join,
2=right join
3=full join

Merge join
has notation
of
JOIN
Or OTRJ

Merg joins
can be used
for inner join
or outer joins

JTAG: if
ends in
1= left join,
2=right join
3=full join

equality check info passed up

Coalesced Var.

```

*****Example 22 *** CORRELATED QUERY *****;
* showing how CORRELATED QUERY is processed ;
proc sql _METHOD _TREE;  TITLE "EX25 A SIMPLE CORRELATED QUERY";
select * from sashelp.class as Outer
Where Outer.AGE = (select Max(age) from sashelp.class as inner where outer.sex=inner.sex);

```

Simple Correlated Query

NOTE: SQL execution methods chosen are:

Sqxsclt (1) this indicates a selection of observations
 Sqxfil (2) this indicates the application of a ?????
 sqxsrlc(SASHELP.CLASS(alias = OUTER)) (20) indicates a selection of observations

NOTE: SQL subquery execution methods chosen are:

Sqxsuqb (1) this indicates a selection of observations
 Sqxsumn (6) this indicates summation without grouping-a summary of the whole table
 sqxsrlc(SASHELP.CLASS(alias = INNER)) (20) indicates a selection of observations

Subquery

Tree as planned.

```

/-SYM-V-(Outer.Name:1 flag=0001)
/-OBJ----| (6)
(3) |--SYM-V-(Outer.Sex:2 flag=0001)
|--SYM-V-(Outer.Age:3 flag=0001)
|--SYM-V-(Outer.Height:4 flag=0001)
\--SYM-V-(Outer.Weight:5 flag=0001)

```

Select * from outer subject to the condition that outer.age = the max age for that sex.

```

/-FIL----| (2)
/-SYM-V-(Outer.Age:3 flag=0001)
/-OBJ----| (11)
(7) |--SYM-V-(Outer.Sex:2 flag=0001)
|--SYM-V-(Outer.Name:1 flag=0001)
|--SYM-V-(Outer.Height:4 flag=0001)
\--SYM-V-(Outer.Weight:5 flag=0001)

```

Pass to FIL

Get this from sashelp.class

```

--SRC----| (4)
\--TABL[SASHELP].class opt=''
(8)
/-SYM-V-(Outer.Age:3)
--CEQ----| (9)
(5)

```

Outer.age is part of outer query & used in the equality check

Fil is the late application of a predicate

Subquery causes the creation of a temp, indexed table

SUBC is short for SUBroutine Call. Printed out as a summary of the process

```

/-SYM-G-(#TEMG001:1 stat=5,0 from Age(0,0) flag=0001)
/-OBJ----| (19)
/-AGGR---| (14)
(12)
|--SYM-V-(inner.Age:3 flag=0001)
|--OBJ----| (25)
|--SRC----| (20)
(15) |--TABL[SASHELP].class opt=''
|--empty-
(21) /-SUBP(1)
\--CEQ----| (26)
(22) \--NAME--(Sex:2)
(27)
|--empty-
|--empty-
|--empty-
|--empty-
|--empty-
(16) /-SYM-G-(#TEMG001:1 stat=5,0 from Age(0,0))
--TLST---| (23)
(17) /-SYM-S-(Age:2 ss=0008x)
\--SLST---| (24)
(18)
\--SUBC---| (10)
(13)

```

SUBP: stands for subroutine parameter

Create a variable, using Grouping (SYM-G), containing the max age (stat function=5)

Pretend that what is inside the paren, in the query, is a function and we are passing , to the function, a parameter- we pass outer.sex from the outer query to the function.

*****Example 23 *** calculated *****;

```
15 proc sql _method _tree;
16 select name, age*12 as age_mo
17 from sashelp.class
18 where calculated age_mo LE 144;
```

NOTE: SQL execution methods chosen are:

```
sqxslct
sqxfil
sqxsrc( SASHELP.CLASS )
```

Tree as planned.

```

      /-SYM-V-(class.Name:1 flag=0001)
    /-OBJ----| (6)
    | (3) \-SYM-A-(age_mo:1 flag=0031)
  /-FIL----|
  | (2) |
  | | /-SYM-V-(class.Age:3 flag=0001)
  | | /-OBJ----| (11)
  | | | (7) \-SYM-V-(class.Name:1 flag=0001)
  | | |--SRC----|
  | | | (4) \-TABL[SASHELP].class opt=''
  | | | /-SYM-A-(age_mo:1 flag=0070)
  | | |--CLE----| (8)
  | | | (5) \-LITN(144)
  | | |--empty- (9)
  | | |--empty-
  | | | /-SYM-A-(age_mo:1 flag=0070)
  | | | /-ASGN---| (12)
  | | | | (10) | /-SYM-V-(class.Age:3)
  | | | | \-AMUL---| (14)
  | | | | (13) \-LITN(12)
  | | \-ERLY---|
  | \-SSEL---| (6)
--SSEL---| (1)
```

Call FIL subroutine in SQL (2) because the variable age_mo was not available to the data engine. It must be created by SQL and then logic can be applied.

"ERLY" is short for "early list."

It contains a list of expressions that must be evaluated before any other expressions on a step.

On that list the first one must be evaluated before the second one (if any), the second one must be evaluated before the third one (if any), etc.

For example, "age_mo" must be calculated (it's on the early list) BEFORE its value can be referenced in the "age_mo <= 144" expression.