

## Increase Your Productivity by Doing Less

Arthur S. Tabachneck, Ph.D., myQNA, Inc., Thornhill, Ontario Canada  
Xia Ke Shan, Chinese Financial Electrical Company, Beijing, China  
Robert Virgile, Robert Virgile Associates, Inc., Lexington, MA  
Joe Whitehurst, High Impact Technologies, Atlanta GA

### ABSTRACT

Using a keep dataset option when declaring a data option has mixed results with various SAS<sup>®</sup> procedures. It might have no observable effect when running PROC MEANS or PROC FREQ but, if your datasets have many variables, it could drastically reduce the time required to run some procs like PROC SORT and PROC TRANSPOSE. This paper describes a fairly simple macro that could easily be modified to use with any proc that defines which variables should be kept and, as a result, make your programs run 12 to 15 times faster.

### THE PROBLEM

In a tutorial presented at the 1996 SUGI meeting, Bruce Gilson presented a number of things SAS programmers can do to improve the efficiency of their programs. One of those methods was to only keep the variables which will be necessary for a particular analysis. Unfortunately, the efficiency was only mentioned as the sixth out of nine methods, and was only mentioned without showing how effective the efficiency could be.

Have you ever considered how much time you could save by excluding irrelevant data from your tasks and analyses? Even with the present paper's authors' over 100 years combined experience as SAS users, and the fact that we are all quite familiar with the technique, most of us seldomly use it. Additionally, many of us have probably wasted time running a procedure that required a sorted dataset, only to discover either that we had failed to sort the incoming data, or that the data weren't sorted because the dataset's sort flag had contained an incorrect value.

**How much time is saved by dropping irrelevant variables?** Presume that you had to run a PROC TRANSPOSE on a dataset like the one that would be created by the following code:

```
data have (drop=months i);  
  array var(*) var1-var1000;  
  do idnum=1 to 10000;  
    date="01dec2010"d;  
    do months=3 to 12 by 3;  
      date=intnx('month',date,3);  
      do i=1 to 1000;  
        var(i) = ceil( 9*ranuni(123) );  
      end;  
      output;  
    end;  
  end;  
run;
```

The above code creates a 40,000 record dataset that has 1,002 variables, namely *IDNUM*, *date*, and 1,000 variables labeled *var1* thru *var1000*. Given that dataset, assume that we want to transpose it to produce a file like the one shown in Example 1 on the following page; i.e., with one record for each of the 10,000 *IDNUMs*, and each record containing four variables that reflect the values of *var1* for each of the four quarters (we didn't think you wanted to see the results for all 10,000 idnums so we only showed the first two).

### Example 1

idnum	var1_Qtr1	var1_Qtr2	var1_Qtr3	var1_Qtr4
1	7	8	3	7
2	6	6	5	4

Many SAS users would use code like the following in order to accomplish the task:

```
proc sort data=have out=need;
  by idnum date;
run;

proc transpose data=need out=want (drop=_) prefix=var1_Qtr;
  by idnum;
  id date;
  var var1;
  format date Qtr1.;
run;
```

The only variables that PROC TRANSPOSE uses are those that are specified in the *by*, *id*, *var* and *copy* statements. Dropping any unnecessary variables when running PROC SORT will result in dramatically improving the performance of your run. And, while not quite as dramatic, it will also have a significant performance improvement effect on your running PROC TRANSPOSE.

The following code would have run significantly faster:

```
proc sort data=have (keep=idnum date var1) out=need noequals;
  by idnum date;
run;

proc transpose data=need out=want (drop=_) prefix=var1_Qtr;
  by idnum;
  id date;
  var var1;
  format date Qtr1.;
run;
```

We tested the differences between the two methods by running 50 trials of each of the two conditions, randomly selecting which method would be run on each trial. Bruce Gilson's efficiency tip allowed PROC SORT to run 7 times faster and PROC TRANSPOSE to run 4.67 times faster.

When we originally wrote this paper for SGF 2013, we only reported the performance gains in terms of the differences in CPU times that we observed, as there weren't any substantial differences between CPU and real times resulting from any of our tests.

Clearly, the originally observed performance differences in themselves provide sufficient justification to always use the efficiency technique when using either procedure. However, in reviewing the paper for the 2013 MWSUG meeting, we ran our code on a different computer and noticed differences between the CPU and real times that had elapsed. As a result, we ran new tests, utilizing a more rigorous testing procedure, namely running 50 trials for each condition, running the trials in a totally random order, comparing both real and cpu times, and running the tests on datasets ranging in size between 40,000 and 640,000 records.

The intent of our tests was to be able to provide definitive benchmarks, but it turned out that too many unforeseeable factors are involved to make that possible (e.g., the number of records and variables being analyzed, the number of variables being dropped, the degree to which the computer was optimized to run SAS, and the number and type of other processes simultaneously competing for the computer or server's resources).

That said, the originally reported performance enhancements appear to have been good estimates, but the expected actual real time differences can only be reliably estimated on a case-by-case basis. What we can say is that **the efficiency optimized sort and transpose processes always ran faster than the non-optimized processes** and, in the worst case we observed, **the efficiency optimized sort process ran more almost 200 times faster than the non-optimized process**.

Thus, the problem addressed by the present paper was whether a method could be offered that would provide users with the best chance of achieving such performance gains and do it in a manner that required less work by both the user and the user's computer.

## THE SOLUTION

Since procedures like PROC TRANSPOSE inherently define all of the variables that will be required, we wrote a fairly simple SAS macro that lets you indicate whether the data has to be sorted and, regardless, ensures that only the relevant variables are used throughout the process. The following call of the macro would produce the desired result, capitalize on the above mentioned time savings, and require you to write less code:

```
%transpose( data=have, out=want (drop=_:), prefix=var1_Qtr, sort=yes,
            by=idnum, var=var1, id=date, format=date Qtr1.)
```

The above call would cause the macro to submit the SAS programs shown, below, as Example 2:

### Example 2 Code that the macro would submit

```
proc sort data=have (keep=idnum date var1) out=_temp noequals;
  by idnum date;
run;

proc transpose data=_temp (keep=idnum date var1) out=want (drop=_:)
  prefix=var1_Qtr;
  by idnum;
  id date;
  var var1;
  format date Qtr1.;
run;

proc delete data=work._temp;
run;
```

We included a *noequals* option in the PROC SORT call because of a suggestion that Philip Mason made in a 2001 SUGI tutorial. Philip suggested that the order of records within by groups seldom has to be maintained, that the order will be maintained unless the option is specified, and that he has typically saved approximately five percent of his processing time by including the option. Since order within by groups is totally irrelevant when one is using PROC TRANSPOSE if the data are sorted by both the by and id variables, we evaluated the option's effect on file sizes of 250,000 and 2.5 million records, when there were three variables and only one by variable. The savings were 12.5 percent on the smaller file and over 14 percent on the larger file.

Philip also suggested that using the *tagSORT* option could save additional processing time, as could using views and compressed datasets, but that such use could also be costly dependent upon the size of one's data and the number of keys specified in a *by* statement. As such, we decided to only include the *noequalS* option, but allow users to specify additional options in a *sort\_options* parameter.

You could achieve the same processing time savings by writing the above code yourself, but the point of the macro is that you necessarily already provide the needed information whenever you run PROC TRANSPOSE. Since such macros are easy to write, and only have to be written once, they result in your having to type less code, lower the chance of your making typographical errors or omissions, and have the added benefit of not requiring you to remember to include the various options. Further, since the macro uses *named parameters*, any settings that you typically use can be set as default values, thus require even less typing.

With the macro, you always get to benefit from the above mentioned processing efficiencies and get to do it by typing less code and having less to remember.

## THE %TRANSPOSE MACRO

The %transpose macro was designed to provide the above mentioned solution. Basically, the program simply creates and runs the code that actually accomplishes the task. The macro's named parameters are the various options and statements that one might use when actually running the procedure, along with two additional parameters that define whether the dataset should be sorted and the options that should be applied. Whether the parameters are actually included in the procedure call is controlled by a series of %if statements.

**Named Parameters.** Named parameters were used so that (1) default values could be assigned and (2) the various parameters would only have to be specified when values other than the default values are required. When calling the macro, the default values will be used unless you specify the desired value. Thus, if you wanted the macro to typically sort your data and result in a file called "work.transposed\_file", you would modify the parameters by specifying them in the macro definition. Example:

```
%macro transpose(data=, out=transposed_file, by=, prefix=, suffix=, label=, let=, name=,
var=, id=, idlabel=, format=, delimiter=, copy=, sort=yes, sort_options= );
```

All of the parameters, except for the *sort* and *sort\_options* parameters, reflect the various options and statements that can be used with PROC TRANSPOSE, and their descriptions can be found in the PROC TRANSPOSE documentation. *Sort* is the parameter you would use to indicate whether the input dataset must first be sorted before the data is transposed. Possible values are YES or NO. If left null, the macro code will set this parameter to equal NO. The *sort\_options* parameter is the parameter you can use to incorporate any additional sort options which you think will reduce your overall processing time (e.g., *tagSORT*, *presorted* and/or *force*).

## CONCLUSION

The purpose of the present paper was to describe and share a SAS macro that could be used to accomplish data transpositions faster and easier than can be accomplished with PROC TRANSPOSE. The principal benefits include: (1) the automatic inclusion of a *keep* dataset option; (2) the ability to sort and transpose in one step; (3) the automatic inclusion of a *noequalS* option for all sorts; and (4) the use of named parameters so that one doesn't have to distinguish between options and statements and frequently used values can be saved in the macro definition.

In addition to satisfying the purpose for which it was designed, the macro can be used as a template for achieving similar benefits with other SAS procedures.

## WHERE TO GET THE MACRO

We did our best to only include carefully written and tested code, but the code may have to be updated from time to time to correct for errors or enhancements that we or others might discover. Additionally, while a copy of the macro is included in this paper, copying and pasting from a pdf file often introduces stylish looking quotation marks which aren't correctly recognized by SAS. As such, we created a page for the paper on sasCommunity.org. The page includes copies of the source code and updated versions of this paper. The page can be found at: [http://www.sascommunity.org/wiki/Increase\\_Your\\_Productivity\\_by\\_Doing\\_Less](http://www.sascommunity.org/wiki/Increase_Your_Productivity_by_Doing_Less)

## DISCLAIMER

The contents of this paper are the work of the authors and do not necessarily represent the opinions, practices or recommendations of the authors' organizations.

## CONTACT INFORMATION

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

Arthur Tabachneck, Ph.D., President  
myQNA, Inc.  
Thornhill, ON Canada  
E-mail: atabachneck@gmail.com

Xia Ke Shan  
Chinese Financial Electrical Company  
Beijing, China  
E-mail: keshan.xia@gmail.com

Robert Virgile  
Robert Virgile Associates, Inc.  
Lexington, MA  
E-mail: rvirgile@verizon.net

Joe Whitehurst  
High Impact Technologies  
Atlanta GA  
E-mail: joewhitehurst@gmail.com

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.

## REFERENCES

*Base SAS® Help and Documentation*, SAS Institute Inc. 2009, Cary, NC, USA,  
<http://support.sas.com/documentation/onlinedoc/base/>

*Configuration and Tuning Guidelines for SAS®9 in Microsoft Windows Server 2008*, Crevar, Margaret, SGF 2011 Proceedings, <http://support.sas.com/resources/papers/WindowsServer2008ConfigurationandTuning.pdf>

*Peeking into Windows to Improve your SAS Performance*, Welch, MP, SCSUG 2011 Proceedings,  
<http://www.scsug.org/SCSUGProceedings/2011/mpwelch1/Peeking%20into%20Windows%20to%20Improve%20Performance.pdf>

*SAS Program Efficiency for Beginners*, Gilson, Bruce, SUGI 1996 Proceedings,  
<http://www.sascommunity.org/sugi/SUGI96/Sugi-96-53%20Gilson.pdf>

*SAS Tips I Learnt While at Oxford*, Mason, Philip, SUGI 2001 Proceedings,  
<http://www2.sas.com/proceedings/sugi26/p020-26.pdf>

*Solving SAS® Performance Problems: Employing Host-Based Tools*, Brown, Tony, SAS Institute Inc., Dallas, TX USA White Paper, <http://support.sas.com/rnd/scalability/papers/practicalperf.pdf>

## APPENDIX I

### THE %TRANPOSE MACRO

```
/**This program runs PROC TRANSPOSE and includes a keep statement to limit the data
that needs to be read and lets you indicate whether the data should be sorted
*
* AUTHORS: Arthur Tabachneck, Xia Ke Shan, Robert Virgile and Joe Whitehurst
* DATE: February 11, 2013

Parameter Descriptions: All of the parameters, except for two, reflect the various
options and statements that can be used with PROC TRANSPOSE and their descriptions
can be found in the PROC TRANSPOSE documentation. The two exceptions are:

*sort* the parameter you would use to indicate whether the input dataset must be
sorted before the data is transposed. Possible values are: YES or NO

*sort_options* The noequals option will be used for all sorts, but this parameter
can be used to specify any other options you want used (e.g., presorted or tagsort)
*/

%macro transpose(data=, out=, by=, prefix=, suffix=, label=, let=, name=, var=,
                id=, idlabel=, format=, delimiter=, copy=, sort=, sort_options= );

/*Populate var parameter in the event it has a null value*/
data _temp;
  set &data. (obs=1 drop=&by. &id. &copy.);
run;

%if %length(&var.) eq 0 %then %do;
  proc sql noprint;
    select name into :var separated by " "
    from dictionary.columns
    where libname="WORK" and memname="_TEMP" and type="num" ;
  quit;
%end;

/*If sort parameter has a value of YES, create a sorted temporary data file*/
%if %sysfunc(upcase("&sort.)) eq "YES" %then %do;
  proc sort data=&data (keep=&by. &id. &var. &copy.) out=_temp
    &sort_options. noequals;
  by &by. &id.;
run;
%let data=_temp;
%end;

/*Run the procedure*/
proc transpose data=&data. (keep=&by. &id. &var. &copy.)
  %if %length(&delimiter.) gt 0 %then delimiter=&delimiter.;
  %if %length(&label.) gt 0 %then label=&label.;
  %if %length(&let.) gt 0 %then let;
  %if %length(&name.) gt 0 %then name=&name.;
  %if %length(&out.) gt 0 %then out=&out.;
  %if %length(&prefix) gt 0 %then prefix=&prefix.;
  %if %length(&suffix) gt 0 %then suffix=&suffix.;;
  %if %length(&by.) gt 0 %then by &by.;;
  %if %length(&copy.) gt 0 %then copy &copy.;;
  %if %length(&id.) gt 0 %then id &id.;;
  %if %length(&idlabel.) gt 0 %then idlabel &idlabel.;;
  %if %length(&var.) gt 0 %then var &var.;;
  %if %length(&format.) gt 0 %then format &format.;;
run;

/*Delete temporary file*/
proc delete data=work._temp; run;
%mend transpose;
```