

# The Focus



A Publication for ANSYS Users

## Contents

### Feature Articles

- [Post Processing: Feasting at the Etable](#)
- [Jobname.solc for Control During a SOLVE](#)
- [Parallel Performance](#)

### On the Web

- [New Solid-Shell Element SOLSH190](#)
- [Materials Database for Purchase](#)

### Resources

- [PADT Support: How can we help?](#)
- [Upcoming Training at PADT](#)
- [About \*The Focus\*](#)
  - *The Focus* Library
  - Contributor Information
  - Subscribe / Unsubscribe
  - Legal Disclaimer

*etables*  
jobname.solc  
Parallel Processing

**ANSYS**  
Support Distributor

# The Focus



A Publication for ANSYS Users

## Post Processing Tips & Hints: Feasting at the Etable

by Eric Miller, PADT

This series on Post Processing with ANSYS has taken a little vacation. Other topics have come up and pushed it aside. But, our erstwhile editor has done a good job to keep reminding us that we need to finish what we started. In this article, we will be looking at one of those very powerful modules that ANSYS has had for a long time and that many users are not all that aware of: E-Tables.

Quoting from the online documentation:

*“Think of the element table as a spreadsheet, where each row represents an element, and each column represents a particular data item for the elements.”*

Back before the power of vector APDL operations in ANSYS, this was the way you extracted and manipulated result data as a group. Also, some types of results information, especially for line elements, are only available by making an etable. Often you will use ETABLEs when the results available using the PRNSOL, PLNSOL do not include the information you need. If that information is stored for the element, you can make an ETABLE, and the plot, list, and otherwise work with this array. If you do anything other than modeling with solid elements, such as composites, beams, coupled-matrix elements, then etable has a lot to offer in terms of getting more information out of your models.

### Do I need to go to a furniture store to get an ETABLE?

One common misconception is that each element has an etable and you call it up to “see” all the other results missing from PRNSOL, PLNSOL, etc. Actually, etables are made on the fly in POST1 and the user tells ANSYS what to fill the rows (elements) and columns (result information) with. This is done with the aptly name ETABLE command and through element selection. You pick what elements you want information for, and then issue an ETABLE command that describes what you want to store and what to call it. ANSYS reads through the results, extracts the information, and places it in the etable. You can make many etables for the same element groups as needed. The command has the format of:

ETABLE, *Lab*, *Item*, *Comp*

The “*Lab*” is the name you want to give the new column, limited to 8 characters (old command), or it can be a command used to (1) refill the table after a SET with REFL, (2) list the status of the table with STAT, or (3) erase the contents of the table with ERAS. *Item* and *Comp* are the same as with most POST1 commands, the *Item* refers to some sort of result type you desire and *Comp* gives a more specific definition (example: stress in x direction is S,X). This command generates a single etable and it holds up

# The Focus



A Publication for ANSYS Users

to 200 columns. If you want to remove a column you can do that by specifying the column name with *Lab* and then use ERAS for the *Item*.

## What can I put on my new Etable?

The information available for placement in an etable varies from element type to element type. If you have ever read the documentation for a given element and actually scrolled to the bottom of the page, you will have noticed some large tables. Open up the ANSYS help and look up SOLID45. If you scroll down to Table 45.1 and 45.2 you will see the typical tables that list “Element Output.” Not everything in these table can be put in an Etable.

The way you can tell if something is available is if a “:” is in the name. The text in front of the colon is the *Item* and the text list after contains all the various *Items* that can be used. Sometimes, as with the Volume of an element, there is no *Comp* argument. If you look at Figure 1, you can see that ETABLE, Vol, VOLUME will store the volume of all the selected elements in the next available column in an ETABLE. ETABLE, StressX, S, X will put the element stresses in the X direction in a column labeled STRESSX.

**Table 45.1 SOLID45 Element Output Definitions**

Name	Definition	O	R
EL	Element Number	Y	Y
NODES	Nodes - I, J, K, L, M, N, O, P	Y	Y
MAT	Material number	Y	Y
VOLU:	Volume	Y	Y
XC, YC, ZC	Location where results are reported	Y	<u>3</u>
PRES	Pressures P1 at nodes J, I, L, K; P2 at I, J, N, M; P3 at J, K, O, N; P4 at K, L, P, O; P5 at L, I, M, P; P6 at M, N, O, P	Y	Y
TEMP	Temperatures T(I), T(J), T(K), T(L), T(M), T(N), T(O), T(P)	Y	Y
FLUEN	Fluences FL(I), FL(J), FL(K), FL(L), FL(M), FL(N), FL(O), FL(P)	Y	Y
S:X, Y, Z, XY, YZ, XZ	Stresses	Y	Y
S:1, 2, 3	Principal stresses	Y	Y
S:INT	Stress intensity	Y	Y



A Publication for ANSYS Users

Figure 1. Typical element output table.

You should note that the R column in the Element Output table tells you which items are available in the .rst file. A Y says it is always there. If there is a number, scroll down to the footnotes under the table. Usually this tells you that you have to set up a special KEYOPT to get the value listed.

If you scroll down further to Table 45.3, you are digging deeper through the strata of ANSYS history to an earlier epoch when short cryptic names and numbers were used to identify results. These are called “Item and Sequence Numbers” and as obtuse as they are, they can be very useful. Since I suspect most of you didn’t actually open up ANSYS help and look at the documentation on SOLID45, you can see table 45.3 shown as Figure 2.

Output Quantity Name	<b>ETABLE and ESOL Command Input</b>								
	Item	I	J	K	L	M	N	O	P
P1	SMISC	2	1	4	3	-	-	-	-
P2	SMISC	5	6	-	-	8	7	-	-
P3	SMISC	-	9	10	-	-	12	11	-
P4	SMISC	-	-	13	14	-	-	16	15
P5	SMISC	18	-	-	17	19	-	-	20
P6	SMISC	-	-	-	-	21	22	23	24
S:1	NMISC	1	6	11	16	21	26	31	36
S:2	NMISC	2	7	12	17	22	27	32	37
S:3	NMISC	3	8	13	18	23	28	33	38
S:INT	NMISC	4	9	14	19	24	29	34	39
S:EQV	NMISC	5	10	15	20	25	30	35	40
FLUEN	NMISC	41	42	43	44	45	46	47	48

Figure 2. Typical Item and Sequence number table.

This table gets a little hard to decipher. To figure out what the Output Quantity Name refers to, look at the previous “Element Output” tables. Then note the label in the Item column. This is what you use as the *Item* in the ETABLE command. For the *Comp* argument, decide what node you want to look at (I-P in this case) and use the appropriate number. So the pressure value on node N on face 3 is obtained with: ETABLE,P3N,SMISC,12. Beam and composite elements can provide a wealth of data using Item and Sequence information.

# The Focus



A Publication for ANSYS Users

## Now that I have an Etable, how do I eat on it?

The primary thing you can do with an etable is plot it on your elements (PLETAB, PLLS) and list it (PRETAB). These commands work like other plot/list commands in ANSYS, but you specify the column — using the name you gave it — instead of the typical *Item, Comp*.

Beyond this ANSYS offers the usual collection of math, sort, and compare operations that you find throughout the program. Table 1 summarizes the various commands you can use to operate on etables. Take a look at the manual pages for each one to get the syntax correct. As you can see, the math options let you do some pretty sophisticated stuff that you may want to use in some advanced post processing. This can be much easier and faster than dumping raw data out to another program.

## Belly up to the Etable.

A quick look at the best place to find ANSYS information on the web, [www.ansys.net](http://www.ansys.net), shows a whole bunch of useful tools that utilize etable commands. Searching [www.xansys.net](http://www.xansys.net) gave more examples of how to use these commands. In certain sections of the ANSYS community, these commands are used every day. So next time you want more information, scroll down the element documentation and look at those tables. What you want is probably there, or can be calculated from it.


Table 1. Etable operation commands

# The Focus



A Publication for ANSYS Users

	Command	Description	Note
Sort	ESORT	Sorts the element table	Use SABS to sort on abs vals
	EUSORT	Restores original order of the element table	
Math	SADD	Adds two element tables	$L_{abR} = (FACT1 \times L_{ab1}) + (FACT2 \times L_{ab2}) + CONST$
	SEXP	Exponent math on an element table	$L_{abR} = ( L_{ab1} ^{EXP1}) \times ( L_{ab2} ^{EXP2})$ Roots, reciprocals, and divides may also be done with this command.
	SMULT	Does multiplication	$L_{abR} = (FACT1 \times L_{ab1}) \times (FACT2 \times L_{ab2})$ Use FACT1 with now Lab2 to scale
	VCROSS	Cross product of two vectors (3 col's per vector)	
	VDOT	Does a dot product with two tables	
Compare	SMAX	Compares two columns and creates a new one with the max	$L_{abR} = (FACT1 \times L_{ab1}) \text{ cmx } (FACT2 \times L_{ab2})$
	SMIN	Compares two columns and creates a new one with the min	$L_{abR} = (FACT1 \times L_{ab1}) \text{ cmn } (FACT2 \times L_{ab2})$
Margin of Safety	SFACT	Does margin of safety calculations	See manual for details
	SFCALC	More safety factor or margin of safety calcs	
	TALLOW	Temp information for factor of safety calculations	
	SALLOW	Sets up allowable stress values for factor of safety plotting.	
Misc	SABS	Specifies absolute values for element table operations	
	SSUM	Sums up tables	Doesn't store anything back in the table, just lists out the values for each column. Use *get,,ssum to store.



A Publication for ANSYS Users

# Jobname.solc for Control During a SOLVE

by Rod Scholl, PADT

Often in the middle of a solution, I wish I had set a parameter differently before hitting SOLVE. Not everything can be changed once the solution's underway, but by creating a jobname.solc you might save yourself from having to start over.

At each equilibrium iteration, ANSYS checks for the existence of a jobname.solc file. If present and modified since the last equilibrium iteration, the contents are read into the ANSYS solver.

Thus when the following jobname.solc is created in the working directory while ANSYS is solving:

A screenshot of a Windows TextPad window. The title bar reads 'TextPad - [C:\My Documents\PADT\Service\Focus\0...'. The menu bar includes 'File', 'Edit', 'Search', 'View', 'Tools', 'Macros', 'Configure', and 'V...'. The active document is named 'ctest.solc'. The text in the window is: 

```
!This is the contents of jobname.solc  
outres,all,all|
```

The output reflects the file was read:

# The Focus



A Publication for ANSYS Users

```

Mark ANSYS 9.0 Output Window
LINE SEARCH PARAMETER = 1.000      SCALED MAX DOF INC = -0.7712E-
  2 CONTACT POINTS HAVE CHANGE OF CONTACT STATUS
FORCE CONVERGENCE VALUE = 115.8      CRITERION= 0.3915E-01
Writing NEWTON-RAPHSON residual forces to file: ctest.nr002
curEqn= 6  totEqn= 1482  Job CP sec= 195.000
Factor Done= 0%  Factor Wall sec= 0.000  rate= 0.0  Mflops
EQUIL ITER 5 COMPLETED.  NEW TRIANG MATRIX.  MAX DOF INC= -0.3114

*** WARNING ***                      CP = 12.469  TIME= 11:11
COMMAND MESSAGE !This is the contents of jobname.solc received from
file: *.SOLC.  Ansys is now processing this Command, but not
responsible for its unexpected side-effects.

WRITE ALL ITEMS TO THE DATABASE WITH A FREQUENCY OF ALL
FOR ALL APPLICABLE ENTITIES

*** WARNING ***                      CP = 12.484  TIME= 11:11
COMMAND MESSAGE outres,all,all received from file: *.SOLC.  Ansys is
now processing this Command, but not responsible for its unexpected
side-effects.

WRITE ALL ITEMS TO THE DATABASE WITH A FREQUENCY OF ALL
FOR ALL APPLICABLE ENTITIES
LINE SEARCH PARAMETER = 1.000      SCALED MAX DOF INC = -0.3114E-
FORCE CONVERGENCE VALUE = 12.29      CRITERION= 0.2801E-01
Writing NEWTON-RAPHSON residual forces to file: ctest.nr003
EQUIL ITER 6 COMPLETED.  NEW TRIANG MATRIX.  MAX DOF INC= -0.1321
  
```

You can use this ability in powerful ways for long-running solutions:

1. Turn on NLDIAG,NRRES if your solution is non-convergent. You can even plot the files in another session (while you pause the running solution, if you want). Perhaps you can affect better convergence by then adding more changes in the jobname.solc.
2. Turn on OUTRES,ALL,ALL. When the solution is divergent, for example, you might want this to grab a converged solution to avoid having only a 999999 unconverged result.
3. Turn on RESCONT, ALL, ALL; if you decide you may need a restart.
4. Turn on NLGEOM!
5. Change Convergence criteria if it is asymptotically approaching your criteria (as seen on the GST) and maybe you can live with the larger residuals.
6. Modify NEQIT to add a few iterations. Nothing worse than when ANSYS bisects, when by the GST you can tell that a couple more iterations would have converged it.
7. Turn on contact debugging tools.
8. **Add DOF constraints!!!!**



# *The Focus*



A Publication for ANSYS Users

On solutions that take a couple days, in another session I often plot .rst at converged substeps and Newton Raphson Residuals via (.nr00x files). With the ability to add DOF constraints between equilibrium iterations, one can be quite creative with getting convergence during a solution, rather than submitting a SOLVE and waiting 2 days just to discover that one more stability DOF was needed.

Just look out for the “unexpected side-effects” warned about. I seem to be losing my hair, but I guess I can’t say ANSYS didn’t warn me.

# The Focus



A Publication for ANSYS Users

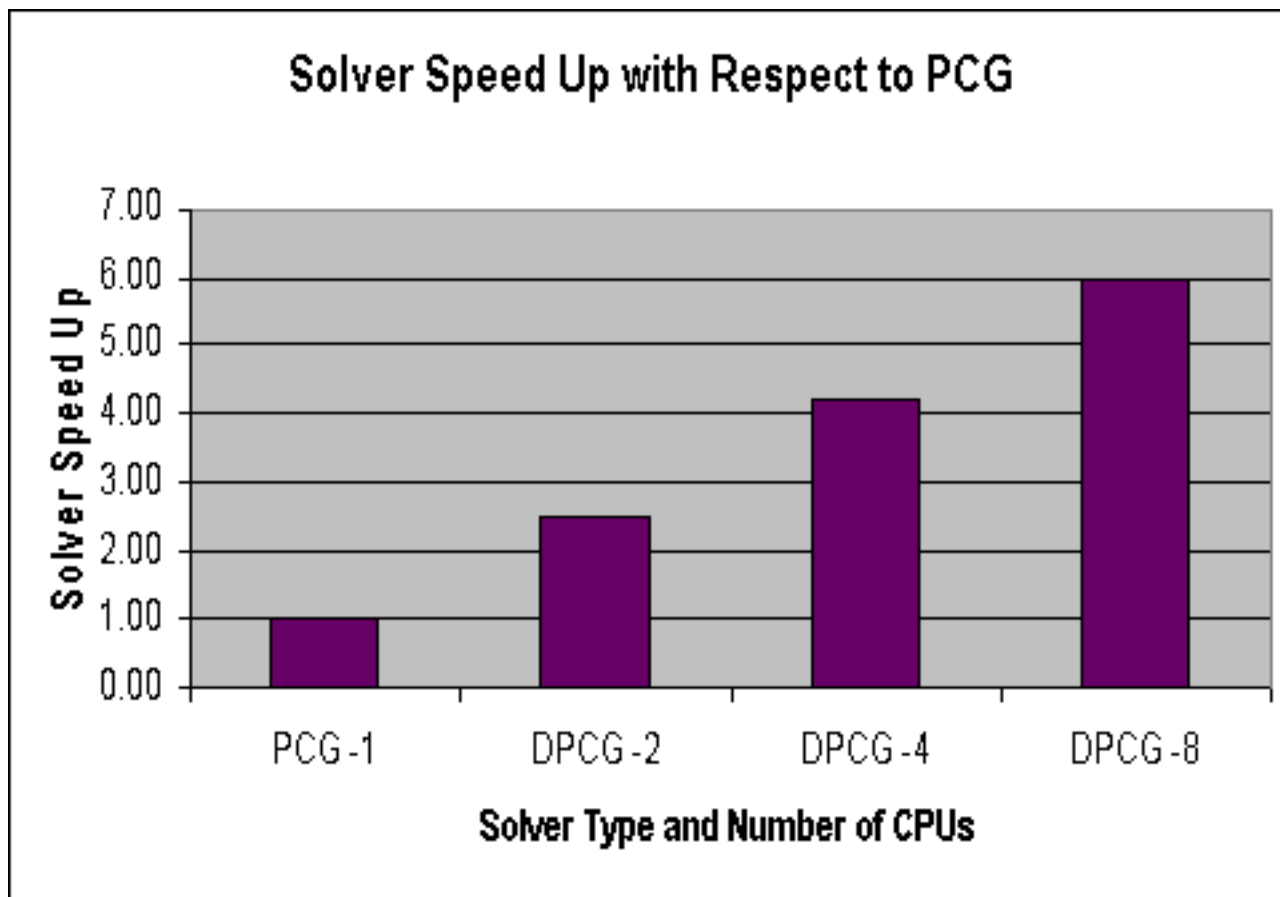
## Parallel Performance

by Stephen Hendry, PADT

Did you ever wish your ANSYS job would run faster and finish sooner? Do you find yourself simplifying your geometry to make your problem size smaller? Do you not have access to Japan's Earth Simulator or Lawrence Livermore National Laboratory's Thunder (first and second fastest computers in the world at 40 TFlops and 22 TFlops, respectively, according to [www.top500.org](http://www.top500.org))? Well then, Parallel Performance for ANSYS could be the answer.

The Parallel Performance for ANSYS Add-on Module facilitates the use of multiprocessor servers, networked workstations, or clusters to generate simulation and analysis results much more quickly.

With two new solvers added at ANSYS 8.0, the Distributed Pre-Conditioned Conjugate Gradient (DPCG) and the Distributed Jacobi Conjugate Gradient (DJCG), even better speed up performance has been achieved. Take, for example, a 12.8 million DOF stress problem with 2.5 million elements. As you can see from the chart below, the DPCG solver has a tremendous advantage as compared to the standard PCG solver.





A Publication for ANSYS Users

## Here is a breakdown of the solvers:

### Distributed Pre-conditioned Conjugate Gradient (DPCG)

- Runs on shared memory or distributed memory
- Superior scalability to PCG solver
- More robust and uses less memory than DDS solver with similar scalability at low (<16) number of processors.
- For an in-depth technical look at using the DPCG solver, take a look at the recent article “[Using the DPCG Solver](#)” by Carlos Shultz in a previous issue of *The Focus*.

### Distributed Jacobi Conjugate Gradient (DJCG)

- Targets distributed parallel processing for single field-type problems, such as thermal analysis.
- Uses both shared and distributed memory architectures.
- Scalability superior to JCG solver with little additional memory required.

### Algebraic Multigrid (AMG)

- Solves ill-conditioned matrices on shared-memory, multiprocessing servers
- Scales well up to eight processors.
- Works well on models where other iterative solvers have difficulty, including ill-conditioned models resulting from large aspect ratios or models with shell and/or beam elements attached to solid elements.

### Distributed Domain Solver (DDS)

- Scalable solver for use in distributed-memory environments or combinations of distributed and shared-memory environments.
- Intended for static or full transient analyses with symmetric matrices.
- DDS divides model into multiple “domains”, each generally 1,000 to 10,000 degrees of freedom in size.

## What is Shared Memory versus Distributed Memory?

Shared Memory is a single shared memory address space accessible by all processors through a common bus architecture. In a Distributed Memory environment, each CPU or computing node has its own memory address space and is not shared. Communication is achieved through a Message Passing Interface (MPI). The DPCG, DJCG, and DDS solvers all support both shared memory and distributed memory architectures. The AMG, however, is only a shared memory solver.



A Publication for ANSYS Users

Parallel Performance for ANSYS works on 32- and 64-bit Linux and Windows clusters as well as a variety of UNIX platforms including HP, IBM, SGI, and Sun. The types of analyses where Parallel Performance can be used include:

- Linear structural
- Nonlinear structural
- Full transient structural
- Steady-state thermal
- Transient thermal

The question you are probably asking right now is how much are these speed increases worth? For a typical Mechanical user, you are going to start seeing reduced wall clock times at just two processors with Parallel Performance. Over time you can very easily equate these reduced wall times into significant savings and better productivity.

# The Focus



A Publication for ANSYS Users

## Upcoming Training at PADT

A complete schedule of the training courses for the next three months is shown below. [Learn more](#) about how the **Training Services** offered by PADT can save you time and money. Or, feel free to drop an e-mail to our training coordinator, [Ted Harris!](#)

Month	Start	End	Course	Description	Location
Jun '05	6/13	6/15	<a href="#">101</a>	<a href="#">Introduction to ANSYS, Part I</a>	Tempe, AZ
	6/20	6/21	<a href="#">201</a>	<a href="#">Basic Structural Nonlinearities</a>	Albuquerque, NM
	6/22	6/23	<a href="#">204</a>	<a href="#">Advanced Contact and Bolt Pretension</a>	Albuquerque, NM
	6/27	6/28	<a href="#">301</a>	<a href="#">Heat Transfer</a>	Irvine, CA
	6/29	6/30	<a href="#">107</a>	<a href="#">ANSYS Workbench DesignModeler</a>	Irvine, CA
Jul '05	7/11	7/13	<a href="#">104</a>	<a href="#">ANSYS Workbench Simulation - Intro</a>	Albuquerque, NM
	7/14	7/15	<a href="#">801</a>	<a href="#">ANSYS Customization with APDL</a>	Albuquerque, NM
	7/25	7/26	<a href="#">203</a>	<a href="#">Dynamics</a>	Albuquerque, NM
	7/27	7/27	<a href="#">108</a>	<a href="#">Paramesh</a>	Albuquerque, NM
Aug '05	8/1	8/3	<a href="#">104</a>	<a href="#">ANSYS Workbench Simulation - Intro</a>	Tempe, AZ
	8/8	8/10	<a href="#">152</a>	<a href="#">AI*Environment</a>	Tempe, AZ
	8/11	8/12	<a href="#">202</a>	<a href="#">Advanced Structural Nonlinearities</a>	Tempe, AZ
	8/15	8/16	<a href="#">501</a>	<a href="#">ANSYS/LS-DYNA</a>	Tempe, AZ
	8/18	8/19	<a href="#">604</a>	<a href="#">Introduction to CFX</a>	Tempe, AZ
	8/24	8/26	<a href="#">402</a>	<a href="#">High Frequency Electromagnetics</a>	Tempe, AZ

Whether it s one of our regularly-scheduled classes at our training facility, or a customized class tailored to your specifications and location, you can be assured that the training you receive will have immediate, positive results on your capabilities in design, analysis, and product and process improvement.





A Publication for ANSYS Users

## About *The Focus*

*The Focus* is a periodic electronic publication published by PADT, aimed at the general ANSYS user. The goal of the feature articles is to inform users of the capabilities ANSYS offers and to provide useful tips and hints on using these products more effectively. *The Focus* may be freely redistributed in its entirety. For administrative questions, please contact [Rod Scholl](#) at PADT.

## *The Focus* Library

All past issues of *The Focus* are maintained in an online [library](#), which can be searched in a variety of different ways.

## Contributor Information

Please don't hesitate to send in a contribution! Articles and information helpful to ANSYS users are very much welcomed and appreciated. We encourage you to send your contributions via e-mail to [Rod Scholl](#).

## Subscribe / Unsubscribe

To subscribe to or unsubscribe from *The Focus*, please visit the PADT e-Publication [subscriptions](#) management page.

## Legal Disclaimer

Phoenix Analysis and Design Technologies (PADT) makes no representations about the suitability of the information contained in these documents and related graphics for any purpose. All such document and related graphics are provided as is without warranty of any kind and are subject to change without notice. The entire risk arising out of their use remains with the recipient. In no event, including inaccurate information, shall PADT be liable for any direct, consequential, incidental, special, punitive or other damages whatsoever (including without limitation, damages for loss of business information), even if PADT has been advised of the possibility of such damages.

The views expressed in *The Focus* are solely those of PADT and are not necessarily those of ANSYS, Inc.