



Haas Computing Services Research Account Orientation

BEST PRACTICES FOR RUNNING APPLICATIONS ON BEAR

General warning

You should not background processes (e.g. Stata, Matlab) via "&", since the Bear's job manager does that for you. The process is still running on the compute node even though it says the job is done, because the only thing that you backgrounded (via '&') was the call to the job manager's **launcher** for the process and not process itself.

Shortly, do not use the '&'; also, use the tools documented in the guidelines to check the status of your job.

Tips for users running Stata and Stata/SE

Stata licenses

If Stata fails to start due to all concurrent licenses being used at the moment, you can check who is using them by issuing the following command from Bear's shell prompt:

bjobs -u all

You can also get more specific information, e.g. to see jobs that were running during the 7 o'clock hour on March 16, issue:

bhist -t -T 03/16/19:

If you are running Stata and would like to see who else is running it, issue

who

from Stata's command prompt.

Stata memory usage model

Stata memory usage is a sum of dataset memory + variables memory

Setting dataset memory

This is done by issuing the command

'set memory <size>'

at the Stata prompt (not Bear's prompt!). **size** is a number that can be followed by a letter '**m**' for megabytes or '**g**' for gigabytes. If no letter is provided, Stata interprets the number as size in kilobytes. Note that once the size is specified, it can not be downsized. So, for example, if you first issued

'set memory 100m'

you would **not** be able release some of it by issuing

'set memory 10m'

although you would be able to allocate more by issuing

'set memory 1000m'

Setting the max number of variables

This is done by issuing the command

'set matsize <numvar>'

where $10 \leq \mathbf{numvar} \leq 800$ in Intercooled Stata, and $10 \leq \mathbf{numvar} \leq 11000$ in Stata/SE.

Setting **numvar** allocates memory in addition to (i.e. on top of) what was requested by the '**set memory**' command.

Running Stata in 'batch mode'

The recommended way to do this is as follows:

stata -b do filename

Stata will execute the commands in filename.do, and will save the output in filename.log.

A more detailed description of how to run Stata in batch mode is available from Stata's website, at <http://www.stata.com/support/faqs/unix/batch.html>

Updating Stata .ado files and add-ons

1.) After logging into bear, run Stata from the prompt using the **full** path, e.g.,

```
[user@bear user]$ /usr/local/stata9/xstata
```

2.) Run any *update* or *install* commands as needed.

3.) Immediately after updates are finished, exit Stata and then use the command 'xstata' or 'xstata-se' **-without the full path-** to get back into Stata as you normally would, e.g.,

```
[user@bear user]$ xstata
```

In all examples above, **xstata** can also be substituted with **xstata-se**, **stata-se**, or **stata**, as desired, but **do not** run Stata with the full path unless installing .ado and other files.

Bear's memory options (for authorized users only)

Some users are authorized to specify the memory for their jobs. The new command format is:

<appname> [memN] [args]

where:

<appname> - required application name, e.g. xstata-se or matlab

[memN] - optional memory requirement, N is number in GB. if omitted, defaults to 8. currently supported values are 8 and 16; unsupported values are silently ignored.

[args] - optional application-specific command-line arguments, up to 8.

For example,

```
[user@bear user]$ stata mem16 -b do testfile
```

For more information

For more research computing information, refer to the following links:

http://groups.haas.berkeley.edu/HCS/research_computing/research.asp

For assistance, please contact the information desk in the Computer Center, or the Helpdesk via:

E-mail - helpdesk@haas.berkeley.edu
Web - <http://helpdesk.haas.berkeley.edu>
Phone - 510/642-0434 or toll-free 866/642-0434