

# Nek5000 Primer

This is a short primer on the basics of using Nek5000. Nek5000 is based on the Nekton 2.0 spectral element code written by Paul Fischer, Lee Ho, and Einar Rønquist in 1986-1991, with technical input from A. Patera and Y. Maday. The Nek5000 development was continued from the 90s through the present by P. Fischer, Jerry Kruse, Julie Mullen, Henry Tufo, and James Lottes.

## 1 Introduction to Nek5000

Nek5000 simulates unsteady incompressible fluid flow with thermal and passive scalar transport. It can handle general two- and three-dimensional domains described by isoparametric quad or hex elements. In addition, it can be used to compute axisymmetric flows. It is a time-stepping based code and does not currently support steady-state solvers, other than steady Stokes and steady heat conduction.

Nek5000 consists of three principal modules: the pre-processor *prex*, the solver *nek5000*, and the post-processor *postx*. *Prex* and *postx* are based upon an X-windows GUI. Recently, *postx* has been extended to output Visual Toolkit (vtk) files, so that output may be viewed in the cave. This is still under development. The solver is written in F77 and C and uses MPI for message passing and some LAPACK routines for eigenvalue computations (depending on the particular solver employed).

In this section, we discuss the input and output files associated with the simulation engine, nek5000.

### 1.1 Input Data

Each simulation is defined by three files, the *.rea* file, the *.usr* file, and the *SIZEu* file. In addition, there is a derived *.map* file that is generated from the *.rea* file by *prex*. Suppose you are doing a calculation called “shear.” The key files defining the problem would be *shear.rea*, *shear.map*, and *shear usr*. *SIZEu* controls (at compile time) the polynomial degree used in the simulation, as well as the space dimension  $d = 2$  or  $3$ .

The file *shear.rea* consists of several sections:

**parameters** These control the runtime parameters such as viscosity, conductivity, number of steps, timestep size, order of the timestepping, frequency of output, iteration tolerances, flow rate, filter strength, etc. There are also a number of free parameters that the user can use as handles to be passed into the user defined routines in the *.usr* file.

**logicals** These determine whether one is computing a steady or unsteady solution, whether advection is turned on, etc.

**geometry** The geometry is specified in an arcane format specifying the *xyz* locations of each of the eight points for each element, or the *xy* locations of each of the four points for each element in 2D. (For several reasons, this format is due to be changed in the future.)

**curvature** This section describes the deformation for elements that are curved.



Figure 1:  $5 \times 3$  mesh generated by *genbox*.

boundary conditions Boundary conditions (BCs) are specified for each face of each element, for each *field* (velocity, temperature, passive scalar #1, etc.). The most common BC is *E*, which indicates that an element is connected to another element. Many of the boundary conditions support either a constant specification or a user defined specification which may be an arbitrary function. For example, a constant Dirichlet BC for velocity is specified by *V*, while a user defined BC is specified by *v*. This upper/lower-case distinction is used for all cases. There are about 70 different types of boundary conditions in all.

restart conditions Here, you can specify a file to use as an initial condition. The initial condition need not be of the same polynomial order as the current simulation. You can also specify that, for example, the velocity come from one file and the temperature from another. The initial time is taken from the restart file, but this can be overridden.

Output specifications Outputs are discussed in a separate section below.

The file *shear usr* would contain set of fortran subroutines that allows the user direct access to all the variables during the course of the simulation. Here, the user may specify spatially varying properties (e.g., viscosity), volumetric heating sources, body forces, and so forth. One can specify arbitrary initial and boundary conditions through the routines *useric()* and *userbc()*. The routine *userchk()* allows the user to interrogate the solution at the end of each timestep for diagnostic purposes. Some such routines are already provided. For example, there are utilities for computing the time average of *u*, *u*<sup>2</sup>, etc. so that one can analyze mean and rms distributions with the postprocessor. There are routine for computing the vorticity inside the solver if one is interested in tracking that, and so forth.

## 2 Getting Started with Meshing

Some basic programs are available to quickly generate spectral element meshes that consist of tensor-product arrays of elements. The generic name for the program is “genbox,” though it has gone through several generations and the current version is *genb6.f*. It works for 2D, axisymmetric, or 3D, and can be used to couple multiple boxes together, so that the overall domain need not be a tensor-product box. In addition, any mesh built by *genbox* can be modified with *prenek*, merged with any other SE mesh by using *prenek*, or modified in the solver through some user-specified mesh deformation. We illustrate some of the basic concepts with the following axisymmetric example.

### 2.1 Uniformly Distributed Mesh

Suppose you wish to simulate flow through an axisymmetric pipe, of radius  $R = 0.5$  and length  $L = 4$ . You estimate that you will need 3 elements in radial (*y*) direction, and 5 in the *x* direction,

as depicted in Fig. 1. This would be specified by the following input file (called *pipe.box*) to genbox:

```
axisymmetric.rea
2                spatial dimension
1                number of fields
#
#  comments:    This is the box immediately behind the
#               refined cylinder in Ugo's cyl+b.l. run.
#
#
#=====
#
Box 1            Pipe
-5 -3           Nelx Nely
0.0  4.0  1.0   x0  x1  ratio
0.0  0.5  1.0   y0  y1  ratio
v ,0 ,A ,W , ,  BC's: (cbx0, cbx1, cby0, cby1, cbz0, cbz1)
```

- The first line of this file supplies the name of an existing 2D .rea file that has the appropriate run parameters (viscosity, timestep size, etc.). These parameters can be modified later, but it is important that axisymmetric.rea be a 2D file, and not a 3D file.
- The second line indicates the number of fields for this simulation, in this case, just 1, corresponding to the velocity field (i.e., no heat transfer).
- The next set of lines just shows how one can place comments into a genbox input file.
- The line that starts with “Box” indicates that a new box is starting, and that the following lines describe a typical box input. Other possible key characters (the first character of Box, “B”) are “C” and “M”, more on those later.
- The first line after “Box” specifies the number of elements in the  $x$  and  $y$  directions. The fact that these values are negative indicates that you want genbox to automatically generate the element distribution along each axis, rather than providing it by hand. (More on this below.)
- The next line specifies the distribution of the 5 elements in the  $x$  direction. The mesh starts at  $x = 0$  and ends at  $x = 4.0$ . The *ratio* indicates the relative size of each element, progressing from left to right. Here,
- The next line specifies the distribution of the 3 elements in the  $y$  direction, starting at  $y = 0$  and going to  $y = 0.5$ . Again, *ratio*=1.0 indicates that the elements will be of uniform height.
- The last line specifies boundary conditions on each of the 4 sides of the box:
  - Lower-case  $v$  indicates that the left ( $x$ ) boundary is to be a velocity boundary condition, with a user-specified distribution determined by routine *userbc* in the .usr file. (Upper-case  $V$  would indicate that the velocity is constant, with values specified in the .rea file.)
  - $O$  indicates that the right ( $x$ ) boundary is an outflow boundary – the flow leaves the domain at the left and the default exit pressure is  $p = 0$ .
  - $A$  indicates that the lower ( $y$ ) boundary is the axis—this condition is mandatory for the axisymmetric case, given the fact that the lower domain boundary is at  $y = 0$ , which corresponds to  $r = 0$ .
  - $W$  indicates that the upper ( $y$ ) boundary is a wall. This would be equivalent to a  $v$  or  $V$  boundary condition, with  $\mathbf{u} = 0$ .

## 2.2 Graded Mesh




Figure 2:  $5 \times 2$  mesh generated by *genbox*.

Suppose you wish to have the mesh be graded, that you have increased resolution near the wall. In this case you change *ratio* in the *y*-specification of the element distribution. For example, changing the 3 lines in the above *genbox* input file from

```
-5 -3          Nelx Nely
0.0  4.0  1.0   x0 x1  ratio
0.0  0.5  1.0   y0 y1  ratio
```

to

```
-5 -4          Nelx Nely
0.0  4.0  1.0   x0 x1  ratio
0.0  0.5  0.7   y0 y1  ratio
```

yields the mesh shown in Fig. 2.

## 2.3 User-Specified Distribution




Figure 3:  $5 \times 4$  mesh generated by *genbox*.

You can also specify your own, precise, distribution of element locations. For example, another graded mesh similar to the one of the preceding example could be built by changing the *genbox* input file to contain:

```
-5 4          Nelx Nely
0.0  4.0  1.0   x0 x1  ratio
0.000  0.250  0.375  0.450  0.500   y0 y1 ... y4
```

Here, the positive number of elements for the *y* direction indicates that *genbox* is expecting *Nely*+1 values of *y* positions on the *y*-element distribution line. This is the *genbox* default, which explains why it corresponds to *Nely* > 0. The corresponding mesh is shown in Fig. 3.

## 2.4 Mesh Modification in Nek5000

For complex shapes, it is often convenient to modify the mesh direction in the simulation code, Nek5000. This can be done through the `usrdat2` routine provided in the `.usr` file. The routine `usrdat2` is called by `nek5000` immediately after the geometry, as specified by the `.rea` file, is established. Thus, one can use the existing geometry to map to a new geometry of interest.

For example, suppose you want the above pipe geometry to have a sinusoidal wall. Let  $\mathbf{x} := (x, y)$  denote the old geometry, and  $\mathbf{x}' := (x', y')$  denote the new geometry. For a domain with  $y \in [0, 0.5]$ , the following function will map the straight pipe geometry to a wavy wall with amplitude  $A$ , wavelength  $\lambda$ :

$$y'(x, y) = y + yA \sin(2\pi x/\lambda).$$

Note that, as  $y \rightarrow 0$ , the perturbation,  $yA \sin(2\pi x/\lambda)$ , goes to zero. So, near the axis, the mesh recovers its original form.

In `nek5000`, you would specify this through `usrdat2` as follows

```
c-----
      subroutine usrdat2
      include 'SIZE'
      include 'TOTAL'
c
      real lambda
c
      ntot = nx1*ny1*nz1*nelt
c
      lambda = 3.
      A      = 0.1
c
      do i=1,ntot
          argx      = 2*pi*xm1(i,1,1,1)/lambda
          ym1(i,1,1,1) = ym1(i,1,1,1) + ym1(i,1,1,1)*A*sin(argx)
      enddo
c
      param(59) = 1. ! Force nek5 to recognize element deformation.
c
      return
      end
c-----
```

Note that, since `nek5000` is modifying the mesh, `postx` will not recognize the current mesh unless you tell it to, because `postx` looks to the `.rea` file for the mesh geometry. The only way for `nek5000` to communicate the new mesh to `postx` is via the `.fld` file, so you must request that the geometry be dumped to the `.fld` file. This is done by modifying the `OUTPUT SPECIFICATIONS`, which are found near the bottom of the `.rea` file. Specifically, change

```
***** OUTPUT FIELD SPECIFICATION *****
6 SPECIFICATIONS FOLLOW
```




Figure 4: Wavy pipe generated by using `usrdat2` to modify the mesh of 3.

```
F      COORDINATES
T      VELOCITY
T      PRESSURE
T      TEMPERATURE
F      TEMPERATURE GRADIENT
O      PASSIVE SCALARS
```

to

```
***** OUTPUT FIELD SPECIFICATION *****
6 SPECIFICATIONS FOLLOW
T      COORDINATES          <----- CHANGE HERE
T      VELOCITY
T      PRESSURE
T      TEMPERATURE
F      TEMPERATURE GRADIENT
O      PASSIVE SCALARS
```

The result of above changes is shown in Fig. 4.

## 2.5 Cylindrical/Cartesian-transition Annuli (update: 7/28/04)

An updated version of `genb6`, known as `genb7`, is currently under development and designed to simply/automate the construction of cylindrical annuli, including *basic* transition-to-Cartesian elements. More sophisticated transition treatments may be generated using the `GLOBAL REFINE` options in `prenek` or through an upgrade of `genb7`, as demand warrants.

Example 2D and 3D input files are provided in the `nek5/doc` files `box7.2d` and `box7.3d`. Figure 5 shows a 2D example generated using the `box7.2d` input file, which reads:

```
x2d.rea
2          spatial dimension
1          number of fields
#
#  comments
#
#
#=====
#
```




Figure 5: *Left*: Cylinder/Cartesian-transition mesh generated by *genb7* using *box7.2d* input file. *Right*: Mesh generated using *box7.2da* input file, followed by one level of quad-refine in prenek.

```

Y                cYlinder
3 -24 1          nelr,nel_theta,nelz
.5 .3            x0,y0 - center of cylinder
ccbb             descriptors: c-cyl, o-oct, b-box (1 character + space)
.5 .55 .7 .8     r0 r1 ... r_nelr
0 1 1           theta0/2pi theta1/2pi ratio
v ,W ,E ,E ,    bc's (3 characters + comma)

```

An example of a mesh built with *box7.2da* is shown in Fig. 5, right. The mesh has been quad-refined once with oct-refine option of prenek. The 3D counterpart to this mesh could be joined to a hemisphere/Cartesian transition built with the spherical mesh option in prenek. (See below.)

## 2.6 Building Spherical Meshes in PRENEK

Spherical meshes can come in many flavors. Three basic ones are the interior of a sphere, the exterior of a sphere, and a spherical shell mesh. All of these can be built with prenek. Here, we give the procedure for the latter two cases, starting with the spherical shell.

### 2.6.1 A Basic Spherical Shell Mesh

The following commands will build a spherical shell mesh, having inner radius 0.5 and outer radius 0.6, using 24 spectral elements.

```
preh
```

```
sph_sh
READ PREVIOUS PARAMETERS
```

```

base_ns                (found in nek5/doc)
BUILD FROM FILE
<cr>                  (take default)

REVIEW/MODIFY

DELETE ELEMENT         (delete single element)

CURVE SIDES
SPHERICAL MESH
0 0 0                 (coordinate of sphere center)
w                     (whole sphere)
n                     (no prolate spheroid)
s                     (spherical shell)
0.5
s                     (spherical shell)
0.6
e                     (exit spherical mesh generator)
BUILD MENU
END   ELEMENTS
ACCEPT MATL,QVOL
SET BCs
SET ENTIRE LEVEL
VELOCITY
FORTRAN FUNCTION      (set bcs to "v")
END   LEVEL
ACCEPT B.C.'s
EXIT

```

The mesh can subsequently be refined with oct- or quad-refine options in GLOBAL refine, upon reentering prenek.

### 2.6.2 A Spherical Shell Mesh in a Cartesian Box

The following commands will build a spherical shell mesh, having inner radius 0.5 and outer radius 0.6, using 24 spectral elements, inside a layered Cartesian box.

```

preh

sph_sh
READ PREVIOUS PARAMETERS
base_ns                (found in nek5/doc)
BUILD FROM FILE
<cr>                  (take default)

REVIEW/MODIFY

DELETE ELEMENT         (delete single element)

```



```

CURVE SIDES
SPHERICAL MESH
0 0 0          (coordinate of sphere center)
w             (whole sphere)
n             (no prolate spheroid)
s             (spherical shell)
0.5
s             (spherical shell)
0.6
c             (extend to Cartesian box)
0.9
c             (extend to Cartesian box)
1.5
c             (extend to Cartesian box)
2.5
c             (extend to Cartesian box)
4.0
c             (extend to Cartesian box)
6.5
c             (extend to Cartesian box)
10.0
e             (exit spherical mesh generator)
BUILD MENU
END   ELEMENTS
ACCEPT MATL,QVOL
SET BCs
SET ENTIRE LEVEL
VELOCITY
FORTRAN FUNCTION      (set bcs to "v")
END   LEVEL
ACCEPT B.C.'s
EXIT

```

The mesh can subsequently be refined with oct- or quad-refine options in GLOBAL refine, upon reentering prenek.

## 2.7 Combining Spherical Shell Meshes with Other Meshes

Note that you can combine a cylinder on, say, the interval  $[-z_0,0]$  with a hemisphere, provided they have the same planform at  $z=0$ . To do so, you would simply build a hemisphere w/ prenek; build a cylinder + transition w/ genb7; start a new prenek session using either of these two .rea files and then IMPORT MESH to bring in the other. Prenek will automatically identify the element interface that exists at  $z=0$  and assign the correct boundary conditions.

In a similar fashion, you can insert a sphere or hemisphere in a tensor-product array of elements by: using genbox to build a tensor product array of elements; use prenek to build a spherical region with matching Cartesian interface; using the tensor-product box to start a new session and importing the spherical region with IMPORT MESH. Prenek will query, "Displace existing elements in box?" If you reply *y* then any existing element that is in the bounding box of the incoming mesh will

be deleted. This provides a convenient way of carving out a space for the new mesh within the originating tensor-product domain.

### 3 README FILES

Below are the contents of the README file in the nek5/src directory. I'm trying to upgrade this to a full-blown latex-based manual.

January 20,2002 pff

- o drive.f now supports constant volume flow rate in x, y, or z direction, depending on whether param54 is 1,2,or 3. x=default for any other value of p54.
- o Global fdm method incorporated as solver and preconditioner for E solve when p116=nelx, p117=nely, p118=nelz. Note that setting one of these to its negative value determines the "primary" direction that is wholly contained within each processor. Thus, if, for example, nelx = -6, nely = 8 and nelz = 5, then the 8x5 array of elements would be partitioned among P processors, and each processor would receive (8x5)/P stacks of depth 6. It is thus relatively important that the product of the number of elements in the remaining secondary and tertiary directions (those not flagged by a minus sign) should be a multiple of P.
- o Dealiasing is currently enabled, whenever p99=2. There is some memory savings to be had if dealiasing is not being used by editing DEALIAS and commenting out the appropriate parameter statements.
- o The comm\_mpi routine has been cleaned up. All vector reductions are performed using mpi\_all\_reduce so, in theory, there should not be a constraint that P=2\*\*cd, provided one isn't using the XXt solver (e.g., if one is using the gfdm solver).
- o The XXt solver is almost in place for the steady conduction case.

-----  
9/10/01, pff

The routines navier5.f and connect1.f have been modified to allow for multiple passive scalars.

---

Nekton consists of three principal modules: the pre-processor, the solver, and the post-processor. The pre- and post-processors, (resp. prenek and postnek) are based upon an X-windows GUI. Recently, postnek has been extended to output Visual Toolkit (vtk) files, so that output may be viewed in the cave. This is still under development.

The solver, nek5000, is written in F77 and C and supports either MPI or Intel's NX message passing libraries. It is a time-stepping based code and does not currently support steady-state solvers, other than steady Stokes and steady heat conduction.

---

To get started, you create a directory, say, nek5, and unpack the tarfile:

```
mkdir nek5
mv tarfile.gz nek5
cd nek5
gunzip tarfile
tar -xvf tarfile
rm tarfile
```

The tarfile will put the source code in the directory src, and will create additional subdirectories: nek5/bin, nek5/2d8, nek5/3d6, and nek5/prep.

nek5/bin contains the shell scripts which I use to start a job. I usually keep these files in my /bin directory at the top level. The scripts will be of use so you can see how I manage the files created by nek5.

nek5/3d6 (2d8) contains an example .rea file, with associated .fld00 file for restarting. Once you "gunzip \*" in all of the directories you should be able to build the executable by typing "makenek" in the 3d6 directory.

For each job that you run, you need a corresponding "subuser.f" file in the nek5/src directory. "subuser.f" provides the user supplied function definitions (e.g., for boundary conditions, initial conditions, etc). I usually keep several subuser.f files in the nek5 directory for different cases under the name of \*.user, e.g., channel.user for channel flow, or cyl.user for flow past a cylinder, etc.

To build the source, simply

```
cd 3d6
makenek (or, "makebk" to make in the background...)
```

To run it, type

```
nekb co1a
```

This will run the job entitled "co1a" in background. The scripts I have set up in nek5/bin are:

```
nek:    runs interactively,
nekl:   runs interactively, but redirects std. out to job.log (& logfile)
nekb:   like nekl, but in the background
nekd:   like nek, but with dbx
```

Each looks for the session name (the prefix of the "\*.rea" file) as an argument.

To build the pre- and post-processors, you will cd to nek5/prep. Type "make" to make postx, and "make -f mpre" to build prex.

If you then put postx and prex into your top-level bin directory you'll be able to invoke the pre and post-processors with the commands "prex" or "postx", etc. (I assume you're already somewhat familiar w/ using pre and post?).

-----

Here's a brief explanation of the nek5000 input format.

First, an overview of the structure of the file:

Section I: Parameters, logical switches, etc.

This section tells nek5000 whether the input file reflects a 2D/3D job, what combination of heat transfer/ Stokes / Navier-Stokes/ steady-unsteady / etc. shall be run.

What are the relevant physical parameters.  
What solution algorithm within Nekton to use, what timestep size or Courant number to use, or whether to run variable DT, etc.

Section II: Mesh Geometry Info

This section gives the number of elements, followed by the

4 (8) vertex pairs (triplets) which specify the corner of each two- (three-) dimensional element.

A subsection which follows specifies which element surfaces are curved. The exact parameter definitions vary according to the type of curved surface. I usually make my own definitions when I need to generate a new curved surface, e.g., spheres or ellipsoids. We use the Gordon-Hall mapping to generate the point distribution within elements, and map according to arc-length for points along the edges of two-dimensional elements. For 3D spheres, I take the element surface point distribution to be the intersection of the great circles generated by the corresponding edge vertices, and again use Gordon-Hall to fill the element interior volume.

The next subsection contains all the boundary condition information for each element. "E" means that the element is connected to the element specified by parameter 1, on the side specified by parameter 2. "T" means a constant temperature along the edge of the element, "t" means a temperature distribution according to a fortran function which the user writes and links in with source code contained in "subuser.f". Similarly, "V" means a constant velocity along the element edge (three components specified in parameters 1, 2, and 3, resp.), and "v" implies a user specified fortran function. "O" is for outflow, "SYM" is for symmetry bc's., etc.

### Section III: Output Info.

This section specifies what data should be output, including field dumps (i.e., X,Y,Z + U,V,W + P + T, or any combination thereof) and time history trace info... e.g., u(t), v(t), etc. at a particular point, for each time step. (Corresponds to a hot wire probe output in some sense.)

Also, if this run is to restart from a previous field dump (e.g., one computed with a lower order, but the same spectral element configuration), that is specified in this section.

This is very brief, but it should give you a road map to the general layout of the input files.

---

### Using PRENEK

This is a brief description of how to modify an existing mesh. Typically it's easiest to modify run parameters just by editing the file. If you wish to modify the geometry, it's (generally) best to do this through prenek. To do this, type "pre" (to begin executing prenek).

Then, click on "READ PREVIOUS PARAMETERS"

Enter "box" with the keyboard.

Then "ACCEPT PARAMETERS"

Then, "BUILD FROM FILE"

Just hit <cr> to accept the default, box.rea.

Then, ACCEPT... ACCEPT... ACCEPT... EXIT

Type a "1" on the keyboard to get a formatted .rea file.

-----  
Ok - hopefully you now have a good .rea file.

You can edit it to change the

viscosity  
number of steps  
Courant number (typ.=2)  
Torder (order of time stepping - typ. 1 in the beginning,  
for more stability, then Torder=2 when you are after  
sensitive results...)  
IOSTEP (frequency of .fld dumps)  
etc.

I can help you with these settings.

=====  
STARTING THE JOB:

If you've compiled your code (with LELT=LELV=sufficiently large), you should now be set to run nek5000. I'll give you some shell scripts, (I keep these in my /bin directory...)

nek: runs interactively,  
nekl: runs interactively, but redirects std. out to job.log (& logfile)  
nekb: like nekl, but in the background

=====  
ENDING THE JOB:

The job will quit after either FINTIME is reached (I never use this) or NSTEPS have completed.

Should you wish to terminate sooner, or get an output *\*right now\**, I've implemented a little check for a file called "ioinfo" into the code.

To get an output at the end of the current step, but keep on running,  
type: "echo 1 > ioinfo".

To get an output at the end of the current step, and then stop the job,  
type: "echo -1 > ioinfo".

To stop the job without dumping the .fld file,  
type: "echo -2 > ioinfo".

---

## CHANGING BOUNDARY CONDITIONS, INITIAL CONDITIONS, ETC.

I've set the genbox2.dat and genbox2.small files to specify fortran boundary conditions at inflow (denoted by lower-case characters)

These are computed in the user accessible source code  
in "../src/subuser.f" (SUBROUTINE USERBC and SUBROUTINE USERIC).

In the subuser.f file which I gave you, I specified  
a blasius profile for the initial and boundary conditions.  
This is the same code I used for my hemisphere runs

### \*\*\*NOTE\*\*\*

You probably will want to change the boundary layer thickness (delta) in  
subuser.f. Just edit subuser.f, search for where delta is specified, change  
it, and go back to 3d6 (3d8? 3d10?) and type "makenek". We can very easily  
make this a runtime parameter, if you find that you're making many trials  
with different boundary layer thicknesses.

---

## RESTARTING

This is really easy...

Just take the .rea file you're about to start with, go to the bottom  
of the file, go up ~33 lines, and change the line:

```
"          0 PRESOLVE/RESTART OPTIONS  *****"
```

to:

```
"          1 PRESOLVE/RESTART OPTIONS  *****"  
"my_old_run"
```

or:

```
"          1 PRESOLVE/RESTART OPTIONS  *****"  
"my_old_run.fld"
```

or:  
"           1 PRESOLVE/RESTART OPTIONS   \*\*\*\*\*"  
"my\_old\_run.fld01   TIME=0.0 "

etc.... (note, drop quotations)

Note that the new run must have the same topology as the old run, that is, the same number of elements in the same location. However, you can change polynomial degree by running in a different directory (e.g., 3d8) with a code which has been compiled with a different SIZEu file.

---

Files needed to start a particular run called, say, "hemi1":

hemi1.rea  
hemi1.sep  
hemi1.map

If hemi1 is to use an old field file (say, "hemi\_old.fld23") as an initial condition, you will also need to have:

hemi\_old.fld23

If hemi\_old.fld23 is in binary format, you will need the associated header file (which is ascii):

hemi\_old.fhd23

Note that hemi\_old must have the same number of elements as hemi1. However, it can be of different polynomial degree, provided it is of degree less than or equal to  $NX1+2$ , where  $NX1$  is the number of grid points in each direction within each element.  
(Note:  $NX1 :=$  polynomial degree + 1.)

In addition to the above hemi1-specific files, you will also need a job specific "subuser.f" file to be compiled and linked in with the source code. Frequently this file is the same over a broad range of parametric studies. It incorporates the user defined forcing functions, boundary conditions, and initial conditions (if not restarting from a previous .fld file). Hence, it's likely that you will not need to edit/recompile this routine very often.

Note that if you are simultaneously undertaking two \*different\* Nekton studies in two different directories which both point to the same source, you will need to swap subuser.f files each time you recompile the source in either directory. I usually do this by keeping the subuser.f files in the src directory under the names, e.g., hemi.user, cylinder.user, channel.user, etc. Then when I'm



compiling the source code for the hemisphere problem I would go to the /src directory and copy hemi.user to subuser.f prior to compiling in the working directory.

---

To build nek5000,

```
cd 3d6
makenek
```

---

To run nek5000 in foreground

```
../bin/nek hemi1
```

To run nek5000 in background

```
../bin/nekb hemi1
```

---

To terminate a job (other than "kill")

```
echo 1 > ioinfo      --- dumps fld file after current step and continues
echo -1 > ioinfo     --- dumps fld file at end of step and quits
echo -2 > ioinfo     --- quits at end of step. No fld file dumped.
echo -# > ioinfo    (#>2) --- quits at end of step #, after dumping fld file.
```

---

To run with multiple passive scalars (e.g., temperature and one or more other convected fields), you should do the following.

- 0) Make sure that you have a version of nek5000 and postx that works with multiple passive scalars. (If you receive these after 9/9/01, you probably do.)
- 1) Assuming you already have a .rea file set for Navier-Stokes plus heat transfer, you can use prex to generate a new .rea/.map/.sep file set that allows for an additional passive scalar. Simply start prenek in the director of interest. Read parameters from your existing .rea file, then select

ALTER PARAMETERS

PASSIVE SCALAR (1)

## WITH CONVECTION (Y)

Hit <cr> through all the parameters (edit these in the .rea file later, with the editor of your choice). Then, read the geometry from your existing file. When you get to the boundary condition menu, you will be prompted for BC's for the new passive scalar. SET ENTIRE LEVEL, Insulated, is a common choice.

- 2) Exit prenek. Note that the default conductivity and rhocp for the new passive scalar is (1.0,1.0). These can be changed by editing the .rea file (if you're not making them functions of time or space). Simply locate the following lines (found right after the parameters)

```
4 Lines of passive scalar data follows2 CONDUCT; 2RHOCp
0.00100      1.00000      1.00000      1.00000      1.00000
1.00000      1.00000      1.00000      1.00000
2.00000      1.00000      1.00000      1.00000      1.00000
1.00000      1.00000      1.00000      1.00000
```

The first 2 lines are the conductivities for the 9 passive scalars. You only need to set the first of these. In the example above, we have `conduct(PS1) =.001`.

The second 2 lines are the rhocp values for the 9 passive scalars. You only need to set the first of these. In the example above, we have `rhocp(PS1) =2.0`.

## 4 Recent Updates

- 01/2002** drive.f now supports constant volume flow rate in x, y, or z direction, depending on whether param54 is 1,2,or 3. x=default for any other value of p54.
- 01/2002** Global fdm method incorporated as solver and preconditioner for E solve when p116=nelx, p117=nely, p118=nelz. Note that setting one of these to its negative value determines the "primary" direction that is wholly contained within each processor. Thus, if, for example, nelx = -6, nely = 8 and nelz = 5, then the 8x5 array of elements would be partitioned among P processors, and each processor would receive (8x5)/P stacks of depth 6. It is thus relatively important that the product of the number of elements in the remaining secondary and tertiary directions (those not flagged by a minus sign) should be a multiple of P.
- 01/2002** Dealiasing is currently enabled, whenever p99=2. There is some memory savings to be had if dealiasing is not being used by editing DEALIAS and commenting out the appropriate parameter statements.
- 01/2002** The comm\_mpi routine has been cleaned up. All vector reductions are performed using mpi\_all\_reduce so, in theory, there should not be a constraint that P=2\*\*cd, provided one isn't using the XXt solver (e.g., if one is using the gfdm solver).

**01/2002** The XXt solver is almost in place for the steady conduction case.

**01/2002** The routines navier5.f and connect1.f have been modified to allow for multiple passive scalars. (9/10/01). 9/10/01, pff