

Summary of Unstructured Mesh Parallel Preprocessor

Scientific Simulations LLC

December 2024

Overview

The `pre_nsu3d` parallel preprocessor performs the following functions:

1. Read input mesh in parallel
2. Construct line data structures for implicit line solver along prismatic cells
3. Partition mesh ensuring no lines are cut across processors
4. Generate coarse agglomerated multigrid levels (AMG)
5. Compute distance function
6. Write out all data in parallel needed by flow solver

Parallel partitioning is done by linking to the ParMETIS or SCOTCH libraries. ParMETIS is a product of the University of Minnesota, and SCOTCH is a product of INRIA in France.

The input formats supported in item 1 are: **cogsg** (VGRID native), **mcell.unf** or ***.mcell** (NSU3D native), **and b8.ugrid** (BIG ENDIAN SOLID MESH stream IO) or **lb8.ugrid** (LITTLE ENDIAN SOLID MESH stream IO). Note that `ugrid` files can be detected as BIG/LITTLE ENDIAN automatically based on the file name extension. The other file endian types must conform to the endian-type compilation of `pre_nsu3d`.

ugrid file formats are preferred due to their widespread use, although `mcell` has traditionally been the NSU3D native format. `Cogsg` and `mcell` files use the fortran sequential access unformatted format. Reading sequential access fortran unformatted files in parallel can be very inefficient with some compilers (notably `ifortV13`). Therefore `pre_nsu3d` reads these files as stream IO, taking into account the record length indicators in the sequential unformatted files. Record lengths larger than 2 GBytes are not supported. For VGRID files, this limit is hit around 125M cells, and for `mcell.unf`, around 500M cells of any single type (i.e. a hybrid mesh of 200M tets and 300M prisms will be readable). The reader will output a message if the record length is larger than the supported value. There is no limit on the size of the `b8.ugrid` files, except if any integer exceeds the values representable with 31 bits (around 2 billion). For VGRID and `mcell` files larger than the limit, a sequential converter code will be supplied to either write these as `b8.ugrid` files or as `mcell8` files, which will also support 64 bit integers for very large meshes. `Tecform`, used for postprocessing, currently supports native `b8.ugrid` and `lb8.ugrid` files as well.

Note that when `cogsg` (VGRID) files are read in, `pre_nsu3d` will automatically perform the prism merging operations in parallel and write out a new `mcell.unf` file which can be used thereafter. Since the resulting `mcell.unf` file has a small number of additional vertices compared to the original VGRID grid, the `dist` and `amg` files will be tagged with this node number and will report an error if read in with the

original cogsg file. Therefore the mcell file should always be used once the merging operation has been performed.

At present, item 4 (AMG levels) is performed on a single computer node and item 5 (distance function) can be time consuming so these are done the first time a new grid is preprocessed, and the results are written out (in parallel) to a file. When the preprocessor is run again on the same grid, if these files exist, they are read in (in parallel) and used directly to generate the final output for the flow solver. Thus, meshes can be repartitioned in a matter of seconds or minutes for most cases once the AMG levels and distance function have been computed and stored to files.

The preprocessor has been tested on grids up to 120 million points (vertices) or 300 million cells (mixed tetrahedral, prisms and pyramids) for up to 32,678 partitions.

The memory requirements of the preprocessor run about 8 Gbytes per 1 million grid points (about 2Gbytes per 1M cells) which is distributed over the parallel machine.

The memory requirements for the AMG level construction are about 700Mbytes per 1 million grid points. Since this portion of the preprocessor executes sequentially, memory requirements may be a bottleneck for large grids. However, for most meshes adequate memory is available on a single multi-core node. As an example, the AMG portion of the preprocessor required 7Gbytes of memory and 82 secs to process a 10 million point mesh on a desktop Linux workstation.

The largest case tried to date was a 120 million point mesh that required 77 GBytes and about 15 minutes to generate the AMG levels on a fat memory cluster node. A fully parallel version of the AMG level generation is under construction.

Three methods are available for distance function calculation: a brute force method, a marching spheres search method, and an Eikonal equation solver. **(The previously supported octree method is currently disabled/not supported.)**

The brute force method takes significant amount of time but scales well on large numbers of cores. This method can be expected to require of the order of 30 minutes for many medium size meshes on moderate numbers of cores. This is the default method as it gives the exact distance function result and is the most robust overall.

The marching spheres method yields the same exact result as the brute force method but runs considerably faster by making use of an efficient search technique. This is enabled through the **drums** library. For large grids, this method is preferable to the brute force search method. The main issue with this approach is that it occasionally results in failure. In this instance, the brute force search method should be used.

The Eikonal equation solver relies on the solution of a PDE to obtain the distance function and scales well, although it requires cpu time equivalent to several hundred flow solver iterations. The Eikonal equation solver method yields an approximate distance function which will not be exactly the same as the other methods.

All methods are setup to write the distance function to a file and to use the file info once it has been generated rather than recalculating the distance function over again.

Configuration Parameters:

pre_nsu3d (r_409 and later) contains several configuration parameters that can be used to control the mode of operation. These can be specified in the input file and are triggered by the header:

OPTIONAL CONFIGURATION PARAMETERS at the end of the file, followed by a list of parameters.

The list of parameters is formatted in a similar manner to a Fortran namelist. The sample input file "input.pre_nsu3d.config" illustrates the use of configuration parameters. The current list includes the following four parameters:

OPTIONAL CONFIGURATION PARAMETERS

```
IPARALLEL_AMG_WRITE = 1
IDIST_METHOD_DEFAULT = 1
#MAX_LINE_SIZE      = 0
MAX_LINE_DIMENSION  = 200
```

Any subset of these parameters may be specified in any order. Misspelt or nonexistent parameter names are ignored (allowing individual lines to be commented out as: #MAX_LINE_SIZE=0 in the above example). The run-time values of the parameters are output to the sdtout file at run time. Thus, any parameters that are not specified are displayed with their default values in stdout. The output file can be used to determine which parameters are available for modification.

Individual parameters are described below:

- *IPARALLEL_AMG_WRITE = 1 (Default): The local amg binary file is written out in parallel for input to the sequential amg3d procedure as a single image binary file. This parallel write has been found to produce write errors on certain hardware configurations, leading to failure of the amg3d procedure. By setting IPARALLEL_AMG_WRITE = 0, the file is output using a sequential write process. This can be used to overcome parallel write errors on the subject hardware.*
- *IDIST_METHOD_DEFAULT = 1 (Default): Sets the default method for calculating the distance function. This is used when the specified *.dist file is not found at run time. (Note that the old method of specifying a number instead of a file name can also be used to select the distance function method). Supported values:*
 - *IDIST_METHOD_DEFAULT = 1: Brute Force method*
 - *IDIST_METHOD_DEFAULT = 2: Octree method : ***No longer supported****
 - *IDIST_METHOD_DEFAULT = 3: Eikonal Equation method (see documentation for additional optional inputs)*
 - *IDIST_METHOD_DEFAULT = 4: Drums Library (faster with same result as Brute Force)*
- *MAX_LINE_SIZE: Determines maximum length of lines allowed. Default = 0, which means no limit. To omit any line creation, specify 0 under the POIN1 FILE NAME entry, as per the original documentation.*

- *MAX_LINE_DIMENSION: Default = 200. For longer lines increase this value. However, large values can lead to excessive memory requirements that may cause failure depending on hardware configuration.*

Sample Timings:

Some sample timings for relevant size meshes are given below.

10 million point mesh divided into 16 partitions on an 8 core workstation (using hyperthreading):

```
-----
TOTAL WALL CLOCK TIME:                0.2056E+03 secs
Time required for reading mesh:        0.2791E+01 secs;  Ratio:  0.01
Time required for partitioning mesh:    0.1277E+03 secs;  Ratio:  0.62
Time required for writing out mesh:     0.7210E+02 secs;  Ratio:  0.35
-----
```

The time to partition is about 2 minutes. Note the write time is rather slow, over 1 minute. This is partly due to the state of the disk and other IO traffic and can vary significantly.

This timing includes reading AMG levels and dist function from files.

In the event these are not available (i.e on the first pass) the following additional resources were required:

Sequential AMG: 7GBytes of RAM, 82 seconds wall clock time

Distance function: 10 minutes of wall clock time

75 million point mesh divided into 2048 partitions on cluster:

```
-----
TOTAL WALL CLOCK TIME:                0.1437E+03 secs
Time required for reading mesh:        0.1100E+02 secs;  Ratio:  0.08
Time required for partitioning mesh:    0.9842E+02 secs;  Ratio:  0.68
Time required for writing out mesh:     0.2419E+02 secs;  Ratio:  0.17
-----
```

In this case the AMG and dist functions when run on the first pass required the following resources:

Sequential AMG: 48 GBytes of RAM, 10 minutes of wall clock time

Distance function: ~2-4 minutes of wall clock time (depending on convergence tolerance)

75M pt mesh 32768 partitions on 16384 cores

```
-----
TOTAL WALL CLOCK TIME:                0.4902E+03 secs
Time required for reading mesh:        0.4862E+02 secs;  Ratio:  0.10
Time required for partitioning mesh:    0.3027E+03 secs;  Ratio:  0.62
Time required for writing out mesh:     0.6369E+02 secs;  Ratio:  0.13
-----
```

Note that the performance degrades somewhat when the number of partitions becomes very large. Also, this case is hyperthreaded (2 threads per core) so there is slowdown per core due to this as well. This was necessary in this case to get the larger number of partitions. Note the IO also degrades, and this is purely a machine specific issue. However, this is a relatively coarse mesh for the large number of partitions requested and finer meshes would likely scale better.

The partitions generated are of high quality and good scalability can be obtained from a CFD solver running on these partitions. As an example, the figure below shows the scalability obtained with NSU3D using up to 32,768 cores on the partitions generated with this preprocessing code. Note that this case contains only 2,200 points per core at 32,768 partitions and better scaling should be achieved with a finer mesh at this level.

