

IOFlexIO

Paul Walker

Date: 2008/05/12 21:03:46

Abstract

Thorn **IOFlexIO** provides I/O methods to output variables in IEEEIO file format. It also implements checkpointing/recovery functionality.

1 Purpose

Thorn **IOFlexIO** uses John Shalf's FlexIO library (see <http://bach.ncsa.uiuc.edu/FlexIO/> for details) to output any type of grid variables (grid scalars, grid functions, and arrays of arbitrary dimension) in the IEEEIO file format.

The thorn registers two I/O methods with the flesh's I/O interface at startup:

- method `IOFlexIOD` outputs all types of grid variables with arbitrary dimensions
- method `IOFlexIO_2D` outputs two-dimensional slices (xy-, xz-, and yz-slice) of three-dimensional grid functions and arrays

Data is written into files named "`<varname>.ieee`" (for method `IOFlexIO`) and "`<varname>_2d_<plane>.ieee`" (for method `IOFlexIO_2D`). Such datafiles can be used for further postprocessing (eg. visualization) or fed back into Cactus via the filereader capabilities of thorn IOUtil.

2 IOFlexIO Parameters

Parameters to control the `IOFlexIO` I/O method are:

- `IOFlexIO::out_every` (steerable)
How often to do periodic `IOFlexIO` output. If this parameter is set in the parameter file, it will override the setting of the shared `IO::out_every` parameter. The output frequency can also be set for individual variables using the `out_every` option in an option string appended to the `IOFlexIO::out_vars` parameter.
- `IOFlexIO::out_vars` (steerable)
The list of variables to output using the **IOFlexIO** I/O method. The variables must be given by their fully qualified variable or group name. The special keyword *all* requests `IOFlexIO` output for all variables. Multiple names must be separated by whitespaces.
An option string can be appended in curly braces to a group/variable name. Supported options are `out_every` (to set the output frequency for individual variables) and hyperslab options (see section 4 for details).
- `IOFlexIO::out_dir`
The directory in which to place the `IOFlexIO` output files. If the directory doesn't exist at startup it will be created.
If this parameter is set to an empty string `IOFlexIO` output will go to the standard output directory as specified in `IO::out_dir`.

3 Serial versus Parallel Output

According to the output mode parameter settings (`IO::out_mode`, `IO::out_unchunked`, `IO::out_proc_every`) of thorn **IOUtil**, thorn **IOFlexIO** will output distributed data either

- in serial into a single unchunked file

```
IO::out_mode      = "onefile"
IO::out_unchunked = "yes"
```

- in parallel, that is, into separate files containing chunks of the individual processors' patches of the distributed array

```
IO::out_mode      = "proc | np"
```

The default is to output data in parallel, in order to get maximum I/O performance. If needed, you can recombine the resulting chunked datafiles into a single unchunked file using the recombining utility program. See section 7 for information how to build the recombining program.

If you have a lot of different variables to recombine you can use the following Bourne shell commands to recombine them. This assumes that the chunked output files for each variable are located in a subdirectory `<varname>_<vardim>d/`. The recombined output file `<varname>.ieee` would then be placed into the current working directory:

```
for var in *_3d ;                               \
do                                               \
{                                               \
    if [ ! -r $var.ieee ] ; then               \
        ieee_recombiner $var/$var.file_0.ieee $var.ieee; \
    fi;                                         \
};                                              \
done
```

4 Output of Hyperslab Data

By default, thorn **IOFlexIO** outputs multidimensional Cactus variables with their full contents resulting in maximum data output. This can be changed for individual variables by specifying a hyperslab as a subset of the data within the N-dimensional volume. Such a subset (called a *hyperslab*) is generally defined as an orthogonal region into the multidimensional dataset, with an origin (lower left corner of the hyperslab), direction vectors (defining the number of hyperslab dimensions and spanning the hyperslab within the N-dimensional grid), an extent (the length of the hyperslab in each of its dimensions), and an optional downsampling factor.

Hyperslab parameters can be set for individual variables using an option string appended to the variables' full names in the `IOFlexIO::out_vars` parameter.

Here is an example which outputs two 3D grid functions `Grid::r` and `Wavetoy::phi`. While the first is output with their full contents at every 5th iteration (overriding the `IOFlexIO::out_every` parameter for this variable), a two-dimensional hyperslab is defined for the second grid function. This hyperslab defines a subvolume to output, starting with a 5 grid points offset into the grid, spanning in the yz-plane, with an extent of 10 and 20 grid points in y- and z-direction respectively. For this hyperslab, only every other grid point will be output.

```
IOFlexIO::out_every = 1
IOFlexIO::out_vars  = "Grid::x{ out_every = 5 }
                    Wavetoy::phi{ origin      = {4 4 4}
                    direction = {0 0 0
```

```

                                0 1 0
                                0 0 1}
extent      = {10 20}
downsample = {2 2}  }"

```

The hyperslab parameters which can be set in an option string are:

- **origin**
This specifies the origin of the hyperslab. It must be given as an array of integer values with N elements. Each value specifies the offset in grid points in this dimension into the N -dimensional volume of the grid variable.
If the origin for a hyperslab is not given, it will default to 0.
- **direction**
The direction vectors specify both the directions in which the hyperslab should be spanned (each vector defines one direction of the hyperslab) and its dimensionality (= the total number of dimension vectors). The direction vectors must be given as a concatenated array of integer values. The direction vectors must not be a linear combination of each other or null vectors.
If the direction vectors for a hyperslab are not given, the hyperslab dimensions will default to N , and its directions are parallel to the underlying grid.
- **extent**
This specifies the extent of the hyperslab in each of its dimensions as a number of grid points. It must be given as an array of integer values with M elements (M being the number of hyperslab dimensions).
If the extent for a hyperslab is not given, it will default to the grid variable's extent. Note that if the origin is set to a non-zero value, you should also set the hyperslab extent otherwise the default extent would possibly exceed the variable's grid extent.
- **downsample**
To select only every so many grid points from the hyperslab you can set the downsample option. It must be given as an array of integer values with M elements (M being the number of hyperslab dimensions).
If the downsample option is not given, it will default to the settings of the general downsampling parameters `IO::out_downsample_[xyz]` as defined by thorn **IOUtil**.

5 Checkpointing & Recovery

Thorn **IOFlexIO** can also be used for creating IEEEIO checkpoint files and recovering from such files later on.

Checkpoint routines are scheduled at several timebins so that you can save the current state of your simulation after the initial data phase, during evolution, or at termination. Checkpointing for thorn **IOFlexIO** is enabled by setting the parameter `IOFlexIO::checkpoint = "yes"`.

A recovery routine is registered with thorn **IOUtil** in order to restart a new simulation from a given **IOFlexIO** checkpoint. The very same recovery mechanism is used to implement a filereader functionality to feed back data into Cactus.

Checkpointing and recovery are controlled by corresponding checkpoint/recovery parameters of thorn **IOUtil** (for a description of these parameters please refer to this thorn's documentation). The parameter `IO::checkpoint_every_walltime.hours` is not (yet) supported.

6 Importing External Data into Cactus with IOFlexIO

In order to import external data into Cactus (eg. to initialize some variable) you first need to convert this data into an IEEEIO datafile which then can be processed by the registered recovery routine of thorn **IOFlexIO**.

The following description explains the IEEEIO file layout of an unchunked datafile which thorn **IOFlexIO** expects in order to restore Cactus variables from it properly. There is also a well-documented example C program provided (`IOFlexIO/doc/CreateIOFlexIOdatafile.c`) which illustrates how to create a datafile with IEEEIO file layout. This working example can be used as a template for building your own data converter program.

1. Actual data is stored as multidimensional datasets in an IEEEIO file.
2. The type of your data as well as its dimensions are already inherited by a dataset itself as meta-information. But this is not enough for **IOFlexIO** to safely match it against a specific Cactus variable. For that reason, the variable's name, its groupname, its grouptype, the timelevel to restore, and the total number of timelevels must be attached to every dataset as attribute information.
3. Finally, the recovery routine needs to know how the datafile to recover from was created:
 - Does the file contain chunked or unchunked data ?
 - How many processors were used to produce the data ?
 - How many I/O processors were used to write the data ?
 - What Cactus version is this datafile compatible with ?

Such information is attached as attributes to the very first dataset in the file. Since we assume unchunked data here the processor information isn't relevant – unchunked data can be fed back into a Cactus simulation running on an arbitrary number of processors.

The Cactus version ID must be present to indicate that grid variables with multiple timelevels should be recovered following the new timelevel scheme (as introduced in Cactus beta 10).

The example C program goes through all of these steps and creates a datafile `x.ieee` in IEEEIO file layout which contains a single dataset named `"grid:x"`, with groupname `"grid::coordinates"`, grouptype `CCTK_GF` (thus identifying the variable as a grid function), the timelevel to restore set to 0, and the total number of timelevels set to 1.

The global attributes are set to `"unchunked" = "yes"`, `nprocs = 1`, and `ioproc_every = 1`.

Once you've built and ran the program you can easily verify if it worked properly with

```
ioinfo -showattrdata x.ieee
```

which lists all objects in the datafile along with their values. Since the single dataset in it only contains zeros it would probably not make much sense to feed this datafile into Cactus for initializing your x coordinate grid function :-)

7 Utility programs provided by IOFlexIO

Thorn **IOFlexIO** provides the following utility programs:

- `ieee_recombiner`
Recombines chunked IEEEIO datafile(s) into a single unchunked IEEEIO datafile.
- `ieee_merge`
Merges the contents of its input files into a single output file.
This might be useful for **IOFlexIO** datafile created by different runs.
- `ieee_extract`
Extracts a hyperslab from all datasets of the input file.
You can select a hyperslab by specifying an origin and an extent (eg. `128x128x128+32+32+32` selects a 128-cubed hyperslab with origin `(32, 32, 32)`).

- `ieee_convert_from_cactus3`
Converts Cactus 3 IEEEIO datafiles into Cactus 4.
It also takes a textfile as input with mapping information for variable names.
- `ioinfo`
Displays the contents of an IEEEIO datafile (number of datasets stored, datatype, rank, dimensions, and number of attributes for each dataset).

All utility programs are located in the `src/util/` subdirectory of thorn **IOFlexIO**. To build the utilities just do a

```
make <configuration>-utils
```

in the Cactus toplevel directory. The executables will then be placed in the `exe/<configuration>/` subdirectory.

All utility programs are self-explaining – just call them without arguments to get a short usage info. If any of these utility programs is called without arguments it will print a usage message.