
Nalu-Wind Documentation

Release 1.2.0

Nalu-Wind Development Team

Mar 21, 2021

CONTENTS

1	User Manual	1
1.1	Building Nalu-Wind	1
1.2	Running Nalu-Wind	13
1.3	Nalu-Wind - Examples	47
2	Developer Manual	57
2.1	Testing Nalu-Wind	57
2.2	Source Code Documentation	59
2.3	Writing Developer Documentation	80
2.4	Writing User Documentation	83
2.5	Building the Documentation	83
2.6	Developer Workflow	84
2.7	Nalu Style Guide	84
2.8	Contributing to Nalu-Wind	85
3	Nalu-Wind - Theory Manual	87
3.1	Low Mach Number Derivation	87
3.2	Supported Equation Set	89
3.3	Discretization Approach	104
3.4	Advection Stabilization	112
3.5	Pressure Stabilization	114
3.6	RTE Stabilization	115
3.7	Nonlinear Stabilization Operator (NSO)	118
3.8	Turbulence Modeling	121
3.9	Supported Boundary Conditions	123
3.10	Overset	138
3.11	Property Evaluations	145
3.12	Coupling Approach	147
3.13	Time discretization	148
3.14	Multi-Physics	149
3.15	Wind Energy Modeling	149
3.16	Topological Support	157
3.17	Adaptivity	157
3.18	Code Abstractions	158
4	Nalu-Wind - Verification Manual	167
4.1	Introduction	167
4.2	2D Unsteady Uniform Property: Convecting Decaying Taylor Vortex	167
4.3	Higher Order 2D Steady Uniform Property: Taylor Vortex	168
4.4	3D Steady Non-isothermal with Buoyancy	173

4.5	3D Steady Non-uniform with Buoyancy	174
4.6	2D Steady Laplace Operator	174
4.7	3D Steady Laplace Operator with Nonconformal Interface	178
4.8	Precursor-based Simulations	178
4.9	Boussinesq Verification	186
4.10	3D Hybrid 1x2x10 Duct: Specified Pressure Drop	189
4.11	3D Hybrid 1x1x1 Cube: Laplace	189
4.12	Fixed Wing Verification Problem	191
4.13	Actuator line simulations coupled to OpenFAST	194
4.14	Open Boundary Condition With Outflow Thermal Stratification	196
4.15	Specified Normal Temperature Gradient Boundary Condition	204
Bibliography		207
Index		211

USER MANUAL

1.1 Building Nalu-Wind

1.1.1 Building Nalu-Wind Semi-Automatically Using Spack

Mac OS X or Linux

The following describes how to build Nalu-Wind and its dependencies mostly automatically on your Mac using [Spack](#). This can also be used as a template to build Nalu-Wind on any Linux system with Spack.

Step 1

This assumes you have a (Homebrew) installation of GCC installed already (we are using GCC 7.3.0). These instructions have been tested on OSX 10.11, MacOS 10.12, and MacOS 10.13. MacOS 10.12/10.13 will not build CMake or Pkg-Config with GCC because they will pick up system header files that have objective C code in them. Therefore we build Nalu-Wind using Spack on MacOS Sierra by using Homebrew to install `cmake` and `pkg-config` and defining these as external packages in Spack (see [packages.yaml](#)).

Step 2

Checkout the official Spack repo from github (we will checkout into `${HOME}`):

```
cd ${HOME} && git clone https://github.com/spack/spack.git
```

Step 3

Add Spack shell support to your `.profile` or `.bashrc` etc, by adding the lines:

```
export SPACK_ROOT=${HOME}/spack
source ${SPACK_ROOT}/share/spack/setup-env.sh
```

Step 4

Run the [setup-spack.sh](#) script from the repo which tries to find out what machine you are on and then copies the corresponding `*.yaml` configuration files to your Spack installation:

```
cd ${HOME} && git clone https://github.com/exawind/build-test.git
cd ${HOME}/build-test/configs && ./setup-spack.sh
```

Step 5

Try `spack info nalu-wind` to see if Spack works. If it does, check the compilers you have available by:

```
machine:~ user$ spack compilers
==> Available compilers
-- clang sierra-x86_64 -----
clang@9.0.0-apple

-- gcc sierra-x86_64 -----
gcc@7.3.0 gcc@6.4.0 gcc@5.5.0
```

Step 6

Install Nalu-Wind with whatever compiler you prefer (it will default to Apple Clang) merely by running `spack install nalu-wind` or by editing and running the `install_nalu_gcc_mac.sh` script from the [build-test](#) repo:

```
cd ${HOME}/build-test/install_scripts && ./install_nalu_gcc_mac.sh
```

That should be it! When using the install script you will see that Spack will install using the constraints we've specified in `shared_constraints.sh` which specifies a much more specific set of Trilinos options for Nalu-Wind that can shorten the build time.

NREL's Eagle Machine

The following describes how to build Nalu-Wind and its dependencies mostly automatically on NREL's Eagle machine using Spack. This can also be used as a template to help build Nalu-Wind on any Linux system with Spack.

Step 1

Login to Eagle, and checkout the <https://github.com/exawind/build-test.git> repo (we will be cloning into the `${HOME}` directory):

```
cd ${HOME} && git clone https://github.com/exawind/build-test.git
```

Step 2

Checkout the official Spack repo from github:

```
cd ${HOME} && git clone https://github.com/spack/spack.git
```

Step 3

Configure your environment in the recommended way. You should purge all modules and load GCC 7.3.0 in your login script. In the example `.bashrc` in the repo we also load Python. If you have problems building with Spack on Eagle, it is most likely your environment has deviated from this recommended one. Even when building with the Intel compiler in Spack, this is the recommended environment at login.

```
module purge
module load gcc/7.3.0
```

Also add Spack shell support to your `.bashrc` as shown in the example `.bashrc` in the repo or the following lines:

```
export SPACK_ROOT=${HOME}/spack
source ${SPACK_ROOT}/share/spack/setup-env.sh
```

Log out and log back in or source your `.bashrc` to get the Spack shell support loaded. Try `spack info nalu-wind` to see if Spack works.

Step 4

Configure Spack for Eagle. This is done by running the `setup-spack.sh` script provided which tries finding what machine you're on and copying the corresponding `*.yaml` file to your Spack directory:

```
cd ${HOME}/build-test/configs && ./setup-spack.sh
```

Step 5

Try `spack info nalu-wind` to see if Spack works.

Step 6

Note the build scripts and `packages.yaml` configuration files provided here adhere to the official versions of the third party libraries we test with, and that you may want to adhere to using them as well. Also note that when you checkout the latest Spack, it also means you will be using the latest packages available if you do not set constraints at install time and the newest packages may not have been tested to build correctly on NREL machines yet. So specifying versions of the TPL dependencies in your `packages.yaml` file for Spack is recommended.

Install Nalu-Wind using a non-GPU login node with the example script `install_nalu_eagle.sh` or edit the script to use the correct allocation and `nice ./install_nalu_eagle.sh`.

That's it! Hopefully the `install_nalu_eagle.sh` script installs the entire set of dependencies and you get a working build of Nalu-Wind on Eagle... after several hours of waiting for it to build.

To build with the Intel compiler, note the necessary change listed in the `install_nalu_eagle.sh` batch script.

To load Nalu-Wind dependencies (you will need Spack's OpenMPI for Nalu-Wind now) into your path you will need to `spack load openmpi %compiler` and `spack load nalu-wind %compiler`, using `%gcc` or `%intel` to specify which to load.

Development Build of Nalu-Wind

When building Nalu-Wind with Spack, Spack will cache downloaded archive files such as *.tar.gz files. However, by default Spack will also erase extracted or checked out ('staged') source files after it has built a package successfully. Therefore if your build succeeds, Spack will have erased the Nalu-Wind source code it checked out from Github.

The recommended way to get a version of Nalu-Wind you can develop in is to checkout Nalu-Wind yourself outside of Spack and build this version using the dependencies Spack has built for you. To do so, checkout Nalu-Wind:

```
git clone https://github.com/exawind/nalu-wind.git
```

Next, create your own directory to build in, or use the existing `build` directory in Nalu-Wind to run the CMake configuration. When running the CMake configuration, point Nalu-Wind to the dependencies by using `spack location -i <package>`. For example in the `build` directory run:

```
cmake -DTrilinos_DIR:PATH=$(spack location -i trilinos) \  
      -DYAML_DIR:PATH=$(spack location -i yaml-cpp) \  
      -DCMAKE_BUILD_TYPE=RELEASE \  
      ..  
make
```

There are also `do-config` scripts available for this according to machine under the `configs` directory [here](#). These scripts may also provide the capability to access and use pre-built dependencies from a set of modules if they are available on the machine. This should allow you to have a build of Nalu-Wind in which you are able to continuously modify the source code and rebuild.

1.1.2 Building Nalu-Wind Manually

Although we recommend installing Nalu-Wind with Spack, if you prefer not to build using Spack, below are instructions which describe the process of building Nalu-Wind by hand. These instructions are an approximation, due to the many differences that can exist across machines.

Linux and OSX

The instructions for Linux and OSX are mostly the same, except on each OS you may be able to use a package manager to install some dependencies for you. Using Homebrew on OSX is one option listed below. Compilers and MPI are expected to be already installed. If they are not, please follow the OpenMPI build instructions. Currently we are recommending OpenMPI v1.10.7 or MPICH 3.3 and GCC v7.3.0. Start by creating a `$(NALU_ROOT_DIR)` to work in.

Homebrew

If using OSX, you can install many dependencies using Homebrew. Install [Homebrew](#) on your local machine and reference the list below for some packages Homebrew can install for you. This allows you to skip the steps describing the build process for each of these applications. You will need to find the location of the applications in which Homebrew has installed them, to use when building Trilinos and Nalu-Wind.

```
brew install openmpi  
brew install cmake  
brew install libxml2  
brew install boost  
brew tap homebrew/science  
brew install superlu43
```


CMake v3.12.4

CMake is provided [here](#).

Prepare:

```
cd ${NALU_ROOT_DIR}/packages
tar xf cmake-3.12.4.tar.gz
```

Build:

```
cd ${NALU_ROOT_DIR}/packages/cmake-3.12.4
./configure --prefix=${NALU_ROOT_DIR}/install/cmake
make
make install
```

SuperLU v4.3

SuperLU is provided [here](#).

Prepare:

```
cd ${NALU_ROOT_DIR}/packages
tar xf superlu_4.3.tar.gz
```

Build:

```
cd ${NALU_ROOT_DIR}/packages/SuperLU_4.3
cp MAKE_INC/make.linux make.inc
```

To find out what the correct platform extension PLAT is:

```
uname -m
```

Edit `make.inc` as shown below (diffs shown from baseline).

```
PLAT          = _x86_64
SuperLUroot   = /your_path/install/SuperLU_4.3 i.e., ${NALU_ROOT_DIR}/install/SuperLU_
↪4.3
BLASLIB       = -L/usr/lib64 -lblas
CC            = mpicc
FORTRAN       = mpif77
```

On some platforms, the `${NALU_ROOT_DIR}` may be mangled. In such cases, you may need to use the entire path to `install/SuperLU_4.3`.

Next, make some new directories:

```
mkdir ${NALU_ROOT_DIR}/install/SuperLU_4.3
mkdir ${NALU_ROOT_DIR}/install/SuperLU_4.3/lib
mkdir ${NALU_ROOT_DIR}/install/SuperLU_4.3/include
cd ${NALU_ROOT_DIR}/packages/SuperLU_4.3
make
cp SRC/*.h ${NALU_ROOT_DIR}/install/SuperLU_4.3/include
```

Libxml2 v2.9.8

Libxml2 is found [here](#).

Prepare:

```
cd ${NALU_ROOT_DIR}/packages
tar -xvf libxml2-2.9.8.tar.gz
```

Build:

```
cd ${NALU_ROOT_DIR}/packages/libxml2-2.9.8
CC=mpicc CXX=mpicxx ./configure --without-python --prefix=${NALU_ROOT_DIR}/install/
↳ libxml2
make
make install
```

Boost v1.68.0

Boost is found [here](#).

Prepare:

```
cd ${NALU_ROOT_DIR}/packages
tar -zxvf boost_1_68_0.tar.gz
```

Build:

```
cd ${NALU_ROOT_DIR}/packages/boost_1_68_0
./bootstrap.sh --prefix=${NALU_ROOT_DIR}/install/boost --with-libraries=signals,regex,
↳ filesystem,system,mpi,serialization,thread,program_options,exception
```

Next, edit `project-config.jam` and add a ‘using mpi’, e.g,

using mpi: /path/to/mpi/openmpi/bin/mpicc

```
./b2 -j 4 2>&1 | tee boost_build_one
./b2 -j 4 install 2>&1 | tee boost_build_intall
```

YAML-CPP 0.6.2

YAML is provided [here](#). Versions of Nalu before v1.1.0 used earlier versions of YAML-CPP. For brevity only the latest build instructions are discussed and the history of the Nalu-Wind git repo can be used to find older installation instructions if required.

Prepare:

```
cd ${NALU_ROOT_DIR}/packages
git clone https://github.com/jbeder/yaml-cpp
cd yaml-cpp && git checkout yaml-cpp-0.6.2
```

Build:

```
cd ${NALU_ROOT_DIR}/packages/yaml-cpp
mkdir build
cd build
cmake -DCMAKE_CXX_COMPILER=mpicxx -DCMAKE_CXX_FLAGS=-std=c++11 -DCMAKE_CC_
↳COMPILER=mpicc -DCMAKE_INSTALL_PREFIX=${NALU_ROOT_DIR}/install/yaml-cpp ..
make
make install
```

Zlib v1.2.11

Zlib is provided [here](#).

Prepare:

```
cd ${NALU_ROOT_DIR}/packages
tar -zxvf zlib-1.2.11.tar.gz
```

Build:

```
cd ${NALU_ROOT_DIR}/packages/zlib-1.2.11
CC=gcc CXX=g++ CFLAGS=-O3 CXXFLAGS=-O3 ./configure --prefix=${NALU_ROOT_DIR}/install/
↳zlib
make
make install
```

HDF5 v1.10.4

HDF5 1.10.4 is provided [here](#).

Prepare:

```
cd ${NALU_ROOT_DIR}/packages
tar -zxvf hdf5-1.10.4.tar.gz
```

Build:

```
cd ${NALU_ROOT_DIR}/packages/hdf5-1.10.4
./configure CC=mpicc FC=mpif90 CXX=mpicxx CXXFLAGS="-fPIC -O3" CFLAGS="-fPIC -O3"
↳FCFLAGS="-fPIC -O3" --enable-parallel --with-zlib=${NALU_ROOT_DIR}/install/zlib --
↳prefix=${NALU_ROOT_DIR}/install/hdf5
make
make install
make check
```

NetCDF v4.6.1 and Parallel NetCDF v1.8.0

In order to support all aspects of Nalu-Wind's parallel models, this combination of products is required.

Parallel NetCDF v1.8.0

Parallel NetCDF is provided on the [Argon Trac Page](#).

Prepare:

```
cd ${NALU_ROOT_DIR}/packages
tar -zxvf parallel-netcdf-1.8.0.tar.gz
```

Build:

```
cd parallel-netcdf-1.8.0
./configure --prefix=${NALU_ROOT_DIR}/install/parallel-netcdf CC=mpicc FC=mpif90_
↪CXX=mpicxx CFLAGS=-O3 CXXFLAGS=-O3 --disable-fortran
make
make install
```

NetCDF v4.6.1

NetCDF is provided [here](#).

Prepare:

```
cd ${NALU_ROOT_DIR}/packages
tar -zxvf netcdf-4.6.1.tar.gz
```

Build:

```
cd netcdf-4.6.1
./configure --prefix=${NALU_ROOT_DIR}/install/netcdf CC=mpicc FC=mpif90 CXX=mpicxx_
↪CFLAGS="-I${NALU_ROOT_DIR}/install/parallel-netcdf/include -O3" LDFLAGS=-L${NALU_
↪ROOT_DIR}/install/parallel-netcdf/lib --enable-pnetcdf --enable-parallel-tests --
↪enable-netcdf-4 --disable-shared --disable-fsync --disable-cdmremote --disable-dap -
↪-disable-doxygen --disable-v2
make -j 4
make check
make install
```

Trilinos

Trilinos is managed by the [Trilinos](#) project and can be found on Github.

Prepare:

```
cd ${NALU_ROOT_DIR}/packages
git clone https://github.com/trilinos/Trilinos.git
cd ${NALU_ROOT_DIR}/packages/Trilinos
mkdir build
cd build
```

Now create a `do-config-trilinos` script with the following recommended options:

```
#!/bin/bash

# The base directory where mpi is located.
# From here you should be able to find include/mpi.h bin/mpicxx, bin/mpiexec, etc.
MPI_ROOT_DIR=/PathToMPI
NALU_ROOT_DIR=/PathToNaluProjectDir
```

(continues on next page)

(continued from previous page)

```

# Note: Don't forget to set your LD_LIBRARY_PATH to $mpi_base_dir/lib
#       You may also need to add to LD_LIBRARY_PATH the lib directory for the compiler
#       used to create the mpi executables.

# TPLS needed by trilinos, possibly provided by HomeBrew on a Mac
#BOOST_ROOT_DIR=/usr/local/Cellar/boost/1.56.0/include/boost/
#SUPERLU_ROOT_DIR=/usr/local/Cellar/superlu/4.3
BOOST_ROOT_DIR=${NALU_BUILD_DIR}/install/boost
SUPERLU_ROOT_DIR=${NALU_BUILD_DIR}/install/SuperLU_4.3
NETCDF_ROOT_DIR=${NALU_BUILD_DIR}/install/netcdf
HDF5_ROOT_DIR=${NALU_BUILD_DIR}/install/hdf5
PARALLEL_NETCDF_ROOT_DIR=${NALU_BUILD_DIR}/install/parallel-netcdf
ZLIB_ROOT_DIR=${NALU_BUILD_DIR}/install/zlib
TRILINOS_ROOT_DIR=${NALU_BUILD_DIR}/install/trilinos

EXTRA_ARGS=$@

# Cleanup old cache before we configure
# Note: This does not remove files produced by make. Use "make clean" for this.
find . -name "CMakeFiles" -exec rm -rf {} \;
rm -f CMakeCache.txt

cmake \
  -DCMAKE_INSTALL_PREFIX=${TRILINOS_ROOT_DIR} \
  -DCMAKE_BUILD_TYPE:STRING=RELEASE
  -DMPI_USE_COMPILER_WRAPPERS:BOOL=ON
  -DMPI_CXX_COMPILER:FILEPATH=${CXX}
  -DMPI_C_COMPILER:FILEPATH=${CC}
  -DMPI_Fortran_COMPILER:FILEPATH=${FC}
  -DKokkos_ENABLE_DEPRECATED_CODE:BOOL=OFF
  -DTpetra_INST_SERIAL:BOOL=ON
  -DTrilinos_ENABLE_CXX11:BOOL=ON
  -DTrilinos_ENABLE_EXPLICIT_INSTANTIATION:BOOL=ON
  -DTpetra_INST_DOUBLE:BOOL=ON
  -DTpetra_INST_COMPLEX_DOUBLE:BOOL=OFF
  -DTrilinos_ENABLE_TESTS:BOOL=OFF
  -DTrilinos_ENABLE_ALL_OPTIONAL_PACKAGES:BOOL=OFF
  -DTrilinos_ASSERT_MISSING_PACKAGES:BOOL=OFF
  -DTrilinos_ALLOW_NO_PACKAGES:BOOL=OFF
  -DTrilinos_ENABLE_Epetra:BOOL=OFF
  -DTrilinos_ENABLE_Tpetra:BOOL=ON
  -DTrilinos_ENABLE_KokkosKernels:BOOL=ON
  -DTrilinos_ENABLE_ML:BOOL=OFF
  -DTrilinos_ENABLE_MueLu:BOOL=ON
  -DXpetra_ENABLE_Kokkos_Refactor:BOOL=ON
  -DMueLu_ENABLE_Kokkos_Refactor:BOOL=ON
  -DTrilinos_ENABLE_EpetraExt:BOOL=OFF
  -DTrilinos_ENABLE_AztecOO:BOOL=OFF
  -DTrilinos_ENABLE_Belos:BOOL=ON
  -DTrilinos_ENABLE_Ipack2:BOOL=ON
  -DTrilinos_ENABLE_Amesos2:BOOL=ON
  -DTrilinos_ENABLE_Zoltan2:BOOL=ON
  -DTrilinos_ENABLE_Ipack:BOOL=OFF
  -DTrilinos_ENABLE_Amesos:BOOL=OFF
  -DTrilinos_ENABLE_Zoltan:BOOL=ON
  -DTrilinos_ENABLE_STK:BOOL=ON
  -DTrilinos_ENABLE_Gtest:BOOL=ON

```

(continues on next page)

(continued from previous page)

```

-DTrilinos_ENABLE_SEACASExodus:BOOL=ON
-DTrilinos_ENABLE_SEACASEpu:BOOL=ON
-DTrilinos_ENABLE_SEACASExodiff:BOOL=ON
-DTrilinos_ENABLE_SEACASNemspread:BOOL=ON
-DTrilinos_ENABLE_SEACASNemslice:BOOL=ON
-DTrilinos_ENABLE_SEACASloss:BOOL=ON
-DTPL_ENABLE_MPI:BOOL=ON
-DTPL_ENABLE_Boost:BOOL=ON
-DBoostLib_INCLUDE_DIRS:PATH=${BOOST_ROOT_DIR}/include
-DBoostLib_LIBRARY_DIRS:PATH=${BOOST_ROOT_DIR}/lib
-DBoost_INCLUDE_DIRS:PATH=${BOOST_ROOT_DIR}/include
-DBoost_LIBRARY_DIRS:PATH=${BOOST_ROOT_DIR}/lib
-DTPL_ENABLE_SuperLU:BOOL=ON
-DSuperLU_INCLUDE_DIRS:PATH=${SUPERLU_ROOT_DIR}/include
-DSuperLU_LIBRARY_DIRS:PATH=${SUPERLU_ROOT_DIR}/lib
-DTPL_ENABLE_Netcdf:BOOL=ON
-DNetCDF_ROOT:PATH=${NETCDF_ROOT_DIR}
-DTPL_Netcdf_PARALLEL:BOOL=ON
-DTPL_ENABLE_Pnetcdf:BOOL=ON
-DPNetCDF_ROOT:PATH=${PARALLEL_NETCDF_ROOT_DIR}
-DPnetcdf_INCLUDE_DIRS:PATH=${PARALLEL_NETCDF_ROOT_DIR}/include
-DPnetcdf_LIBRARY_DIRS:PATH=${PARALLEL_NETCDF_ROOT_DIR}/lib
-DTPL_ENABLE_HDF5:BOOL=ON
-DHDF5_ROOT:PATH=${HDF5_ROOT_DIR}
-DHDF5_NO_SYSTEM_PATHS:BOOL=ON
-DTPL_ENABLE_Zlib:BOOL=ON
-DZlib_INCLUDE_DIRS:PATH=${ZLIB_ROOT_DIR}/include
-DZlib_LIBRARY_DIRS:PATH=${ZLIB_ROOT_DIR}/lib
-DTPL_ENABLE_BLAS:BOOL=ON
$EXTRA_ARGS \
../

```

Build

Place into the build directory, the `do-config-trilinos` script created from the recommended Trilinos configuration listed above.

`do-config-trilinos` will be used to run `cmake` to build trilinos correctly for Nalu-Wind.

Make sure all other paths to `netcdf`, `hdf5`, etc., are correct.

```

./do-config-trilinos
make
make install

```

HYPRE

Nalu-Wind can use HYPRE solvers and preconditioners, especially for Pressure Poisson solves. However, this dependency is optional and is not enabled by default. Users wishing to use HYPRE solver and preconditioner combination must compile HYPRE library and link to it when building Nalu-Wind.

```

# 1. Clone hypre sources
https://github.com/LLNL/hypre.git

```

(continues on next page)

(continued from previous page)

```
cd hypre/src

# 2. Configure HYPRE package and pass installation directory
./configure --prefix=${NALU_ROOT_DIR}/install/hypre --without-superlu --without-
↳openmp --enable-bigint

# 3. Compile and install
make && make install
```

Note:

1. Make sure that `--enable-bigint` option is turned on if you intend to run linear systems with > 2 billion rows. Otherwise, `nalu` executable will throw an error at runtime for large problems.
2. Users must pass `-DENABLE_HYPRE` option to CMake during Nalu-Wind configuration phase. Optionally, the variable `-DHYPRE_DIR` can be used to pass the path of HYPRE install location to CMake.

ParaView Catalyst

Optionally enable [ParaView Catalyst](#) for in-situ visualization with Nalu-Wind. These instructions can be skipped if you do not require in-situ visualization with Nalu-Wind. The first thing you will need to do is build Paraview yourself using their SuperBuild instructions.

Build Nalu-Wind ParaView Catalyst Adapter

Next you will need to build the Catalyst adapter for Trilinos to hook into Paraview. The adapter is located in the Trilinos repo at `Trilinos/packages/seacas/libraries/ioss/src/visualization/ParaViewCatalystIossAdapter`. To install:

```
cd Trilinos/packages/seacas/libraries/ioss/src/visualization/
↳ParaViewCatalystIossAdapter
cmake -DParaView_DIR:PATH=/path/to/paraview/lib/cmake/paraview_version -DCMAKE_
↳INSTALL_PREFIX:PATH=${NALU_ROOT_DIR}/install
make
make install
```

Nalu-Wind

Nalu-Wind is provided [here](#). The master branch of Nalu-Wind typically matches with the master branch or develop branch of Trilinos. If it is necessary to build an older version of Nalu-Wind, refer to the history of the Nalu git repo for instructions on doing so.

Prepare:

```
git clone https://github.com/Exawind/nalu-wind.git
```

Build

Create a `nalu-wind/build` directory and execute something similar to following commands. The general commands for configuring and building Nalu-Wind are listed below. We show a script which uses modules which populate the `<PACKAGE>_ROOT_DIR` locations for the NREL Eagle machine, but it will need to be modified with the specific TPL locations you have used.

```
#!/bin/bash -l

# Instructions:
# Make a directory in the Nalu-Wind directory for building,
# Copy this script to that directory and edit the
# options below to your own needs and run it.

CXX_COMPILER=mpicxx
C_COMPILER=mpicc
FORTRAN_COMPILER=mpifort
FLAGS="-O2 -march=native -mtune=native"
OVERSUBSCRIBE_FLAGS="--use-hwthread-cpus --oversubscribe"

set -e

TRILINOS_ROOT_DIR=${NALU_ROOT_DIR}/install/trilinos
YAML_CPP_ROOT_DIR=${NALU_ROOT_DIR}/install/yaml-cpp

# Clean before cmake configure
set +e
rm -rf CMakeFiles
rm -f CMakeCache.txt
set -e

# Extra TPLs that can be included in the cmake configure:
# -DENABLE_PARAVIEW_CATALYST:BOOL=ON \
# -DPARAVIEW_CATALYST_INSTALL_PATH:PATH=${CATALYST_IOSS_ADAPTER_ROOT_DIR} \
# -DENABLE_OPENFAST:BOOL=ON \
# -DOpenFAST_DIR:PATH=${OPENFAST_ROOT_DIR} \
# -DENABLE_HYPRE:BOOL=ON \
# -DHYPRE_DIR:PATH=${HYPRE_ROOT_DIR} \
# -DENABLE_TIOGA:BOOL=ON \
# -DTIOGA_DIR:PATH=${TIOGA_ROOT_DIR} \

(set -x; cmake \
  -DCMAKE_CXX_COMPILER:STRING=${CXX_COMPILER} \
  -DCMAKE_CXX_FLAGS:STRING="${FLAGS}" \
  -DCMAKE_C_COMPILER:STRING=${C_COMPILER} \
  -DCMAKE_C_FLAGS:STRING="${FLAGS}" \
  -DCMAKE_Fortran_COMPILER:STRING=${FORTRAN_COMPILER} \
  -DCMAKE_Fortran_FLAGS:STRING="${FLAGS}" \
  -DMPI_CXX_COMPILER:STRING=${CXX_COMPILER} \
  -DMPI_C_COMPILER:STRING=${C_COMPILER} \
  -DMPI_Fortran_COMPILER:STRING=${FORTRAN_COMPILER} \
  -DMPIEXEC_PREFLAGS:STRING="${OVERSUBSCRIBE_FLAGS}" \
  -DTrilinos_DIR:PATH=${TRILINOS_ROOT_DIR} \
  -DYAML_DIR:PATH=${YAML_CPP_ROOT_DIR} \
  -DCMAKE_BUILD_TYPE:STRING=RELEASE \
  -DENABLE_DOCUMENTATION:BOOL=OFF \
  -DENABLE_TESTS:BOOL=ON \
```

(continues on next page)

(continued from previous page)

```
..)

(set -x; nice make -j 16)
```

This process will create `nalux` within the `nalu-wind/build` location.

1.2 Running Nalu-Wind

This section describes the general process of setting up and executing Nalu-Wind, understanding the various input file options available to the user, and how to extract results and analyze them. For the simplest case, Nalu-Wind requires the user to provide a YAML input file with the options that control the run along with a computational mesh in Exodus-II format. More complex setups might require additional files:

- Trilinos MueLu preconditioner configuration in XML format
- ParaView Cataylst input file for in-situ visualizations
- Additional Exodus-II mesh files for solving different physics equation sets on different meshes, or for solution transfer to an input/output mesh.

1.2.1 Exodus-II File Format

Nalu-Wind requires the user to provide the computational mesh in [Exodus-II](#) format. The output and restart files generated by Nalu-Wind are also in Exodus-II format where the requested fields are output along side the mesh. The restart files from one Nalu-Wind simulation can serve as the input file for a subsequent simulation.

Several commercial mesh generation software support output to Exodus-II format. Two such software used by Nalu-Wind developers are:

- [CUBIT](#)
- [Pointwise](#)

Furthermore, [NaluWindUtils](#) provides an `abl_mesh` utility that can be used to generate simple structured meshes (output into Exodus-II format) for use with atmospheric boundary layer simulations.

Examining Exodus-II Files

Exodus-II uses the [NetCDF](#) format to store data, therefore, the several NetCDF utilities can be used to examine the file metadata. For example, the following code snippet shows the use of `ncdump` to examine the names of the mesh blocks and side sets, as well as the nodal fields available in a given mesh file.

```
ncdump -v eb_names,ss_names,name_nod_var channel_coarse_ic.g
# <output truncated to show only relevant parts>
data:

  eb_names =
    "interior" ;

  ss_names =
    "inlet",
    "outlet",
    "bottomwall",
    "topwall",
```

(continues on next page)

(continued from previous page)

```
"back",  
"front" ;  
  
name_nod_var =  
"turbulent_ke",  
"velocity_x",  
"velocity_y",  
"velocity_z" ;
```

For brevity, the example above has removed the NetCDF `dimensions` and `variables` sections to show just the contents of the variable names of interest. The output shows that the mesh in question contains one element block (interior) with six boundary planes (side-sets) and has two nodal fields: the velocity vector, and the turbulent kinetic energy scalar. `ncdump` can be invoked with the `-h` flag to print just the headers. Of particular interest is the NetCDF `dimensions` section that contains information about the total number of nodes, element, boundary faces, etc. in the mesh file.

Most visualization programs support loading Exodus-II mesh/solution files and can be used to visualize the flow fields generated by Nalu-Wind. Two open-source visualization programs available are:

- [ParaView](#)
- [VisIt](#)

Preliminary support for in-situ visualization using [ParaView Catalyst](#) is available within the Nalu-Wind code base and can be enabled by linking to Catalyst libraries during compile time. See input file specifications more details on setting up Cataylst for in-situ visualization of Nalu-Wind solution files.

Other Exodus-II Utilities

A brief description of some useful Exodus-II utilities are provided here. Please consult the documentation of these programs to understand the full range of options available.

decomp

`decomp` is a SEACAS utility (available from a Trilinos install) that can be used to decompose a mesh file across several MPI ranks for use in a subsequent parallel simulation.

eput

`eput` performs the reverse action of `decomp`, i.e., it combines parallel decomposed files from a simulation into a single Exodus-II database. The simplest invocation is

```
eput -auto nalu_output.e.8.0
```

The `-auto` flag determines the database structured based on the file provided on the command line and combines the files (in the above example into `nalu_output.e`).

mapvar-kd

Map solution fields from one mesh to another mesh.

percept

The [Percept](#) project provides various tools to perform mesh refinement, higher-order promotion, etc. See documentation for `mesh_adapt` to determine various options available.

1.2.2 Invoking Nalu-Wind - Command-line options

Nalu-Wind's runtime behavior can be controlled by using several command line input options during invocation. Users can invoke `-h` to determine the various options available.

- h, --help**
Print the help message describing all Nalu-Wind options and exit
- i, --input-deck**
Use the filename provided as the input file. If this option is not provided, **nalux** will attempt to load a file called `nalux.i` in the current working directory as the input file.
- o, --log-file**
The log file where the output generated by Nalu-Wind is directed to. If no file is provided, then **nalux** will use the base name of the Nalu-Wind input file with the extension `.log` as the output file. For example, if **nalux** was invoked as `nalux -i ABL.neutral.i` then the output will be redirected to a file named `ABL.neutral.log`. Note that the file is overwritten if it already exists.
- v, --version**
Print the Nalu-Wind version string.
- p, --pprint**
Enable parallel printing from all MPI ranks.
- D, --debug**
Enable verbose debug printing to log file.

1.2.3 Nalu-Wind Input File

Nalu-Wind requires the user to provide an input file, in YAML format, during invocation at the command line using the `nalux -i` flag. By default, **nalux** will look for `nalux.i` in the current working directory to determine the mesh file as well as the run setup for execution. A sample `nalux.i` is shown below:

Listing 1.1: Sample Nalu-Wind input file for the Heat Conduction problem

```
# -*- mode: yaml -*-
#
# Example Nalu input file for a heat conduction problem
#
Simulations:
- name: sim1
  time_integrator: ti_1
  optimizer: opt1
linear_solvers:
- name: solve_scalar
  type: tpetra
  method: gmres
  preconditioner: sgs
  tolerance: 1e-3
  max_iterations: 75
  kspace: 75
  output_level: 0
realms:
```

(continues on next page)

(continued from previous page)

```
- name: realm_1
mesh: periodic3d.g
use_edges: no
automatic_decomposition_type: rcb

equation_systems:
  name: theEqSys
  max_iterations: 2

  solver_system_specification:
    temperature: solve_scalar

  systems:
    - HeatConduction:
        name: myHC
        max_iterations: 1
        convergence_tolerance: 1e-5

initial_conditions:
  - constant: ic_1
    target_name: block_1
    value:
      temperature: 10.0

material_properties:
  target_name: block_1
  specifications:
    - name: density
      type: constant
      value: 1.0
    - name: thermal_conductivity
      type: constant
      value: 1.0
    - name: specific_heat
      type: constant
      value: 1.0

boundary_conditions:
  - wall_boundary_condition: bc_left
    target_name: surface_1
    wall_user_data:
      temperature: 20.0

  - wall_boundary_condition: bc_right
    target_name: surface_2
    wall_user_data:
      temperature: 40.0

solution_options:
  name: myOptions

  use_consolidated_solver_algorithm: yes
```

(continues on next page)

(continued from previous page)

```

options:
- element_source_terms:
    temperature: FEM_DIFF

output:
  output_data_base_name: femHC.e
  output_frequency: 10
  output_node_set: no
  output_variables:
    - dual_nodal_volume
    - temperature

Time_Integrators:
- StandardTimeIntegrator:
    name: ti_1
    start_time: 0
    termination_step_count: 25
    time_step: 10.0
    time_stepping_type: fixed
    time_step_count: 0
    second_order_accuracy: no

realms:
  - realm_1

```

Nalu-Wind input file contains the following top-level sections that describe the simulation to be executed.

Realms

Realms describe the computational domain (via mesh input files) and the set of physics equations (Low-Mach Navier-Stokes, Heat Conduction, etc.) that are solved over this particular domain. The list can contain multiple computational domains (*realms*) that use different meshes as well as solve different sets of physics equations and interact via *solution transfer*. This section also contains information regarding the initial and boundary conditions, solution output and restart options, the linear solvers used to solve the linear system of equations, and solution options that govern the discretization of the equation set.

A special case of a realm instance is the input-output realm; this realm type does not solve any physics equations, but instead serves one of the following purposes:

- provide time-varying boundary conditions to one or more boundaries within one or more of the participating realms in the simulations. In this context, it acts as an *input* realm.
- extract a subset of data for output at a different frequency from the other realms. In this context, it acts as an *output* realm.

Inclusion of an input/output realm will require the user to provide the additional *transfers* section in the Nalu-Wind input file that defines the solution fields that are transferred between the realms. See *Physics Realm Options* for detailed documentation on all Realm options.

Linear Solvers

This section configures the solvers and preconditioners used to solve the resulting linear system of equations within Nalu-Wind. The linear system convergence tolerance and other controls are set here and can be used with multiple systems across different realms. See *Linear Solvers* for more details.

Time Integrators

This section configures the time integration scheme used (first/second order in time), the duration of simulation, fixed or adaptive timestepping based on Courant number constraints, etc. Each time integration

section in this list can accept one or more `realms` that are integrated in time using that specific time integration scheme. See [Time Integration Options](#) for complete documentation of all time integration options available in Nalu-Wind.

Transfers

An optional section that defines one or more solution transfer definitions between the participating `realms` during the simulation. Each transfer definition provides a mapping of the to and from realm, part, and the solution field that must be transferred at every timestep during the simulation. See [ABL Forcing](#) section for complete documentation of all transfer options available in Nalu-Wind.

Simulations

Simulations provides the top-level architecture that orchestrates the time-stepping across all the realms and the required equation sets.

Linear Solvers

The `linear_solvers` section contains a list of one or more linear solver settings that specify the solver, preconditioner, convergence tolerance for solving a linear system. Every entry in the YAML list will contain the following entries:

Note: The variable in the `linear_solvers` subsection are prefixed with `linear_solvers.name` but only the variable name after the period should appear in the input file.

`linear_solvers.name`

The key used to refer to the linear solver configuration in `equation_systems.solver_system_specification` section.

`linear_solvers.type`

The type of solver library used.

Type	Description
<code>tpetra</code>	Tpetra data structures and Belos solvers/preconditioners
<code>hypre</code>	Hypr data structures and Hypr solver/preconditioners

`linear_solvers.method`

The solver used for solving the linear system.

When `linear_solvers.type` is `tpetra` the valid options are: `gmres`, `biCgStab`, `cg`. For `hypre` the valid options are `hypre_boomerAMG` and `hypre_gmres`.

Options Common to both Solver Libraries

`linear_solvers.preconditioner`

The type of preconditioner used.

When `linear_solvers.type` is `tpetra` the valid options are `sgs`, `mt_sgs`, `muelu`. For `hypre` the valid options are `boomerAMG` or `none`.

`linear_solvers.tolerance`

The relative tolerance used to determine convergence of the linear system.

`linear_solvers.max_iterations`

Maximum number of linear solver iterations performed.

`linear_solvers.kspace`

The Krylov vector space.

linear_solvers.output_level

Verbosity of output from the linear solver during execution.

linear_solvers.write_matrix_files

A boolean flag indicating whether the matrix, the right hand side, and the solution vector are written to files during execution. The matrix files are written in MatrixMarket format. The default value is `no`.

Additional parameters for Belos Solver/Preconditioners**linear_solvers.muelu_xml_file_name**

Only used when the `linear_solvers.preconditioner` is set to `muelu` and specifies the path to the XML filename that contains various configuration parameters for Trilinos MueLu package.

linear_solvers.recompute_preconditioner

A boolean flag indicating whether preconditioner is recomputed during runs. The default value is `yes`.

linear_solvers.reuse_preconditioner

Boolean flag. Default value is `no`.

linear_solvers.summarize_muelu_timer

Boolean flag indicating whether MueLu timer summary is printed. Default value is `no`.

Additional parameters for Hypre Solver/Preconditioners

The user is referred to [Hypre Reference Manual](#) for full details on the usage of the parameters described briefly below.

The parameters that start with `bamg_` prefix refer to options related to Hypre's BoomerAMG preconditioner.

linear_solvers.bamg_output_level

The level of verbosity of BoomerAMG preconditioner. See `HYPRE_BoomerAMGSetPrintLevel`. Default: 0.

linear_solvers.bamg_coarsen_type

See `HYPRE_BoomerAMGSetCoarsenType`. Default: 6

linear_solvers.bamg_cycle_type

See `HYPRE_BoomerAMGSetCycleType`. Default: 1

linear_solvers.bamg_relax_type

See `HYPRE_BoomerAMGSetRelaxType`. Default: 6

linear_solvers.bamg_relax_order

See `HYPRE_BoomerAMGSetRelaxOrder`. Default: 1

linear_solvers.bamg_num_sweeps

See `HYPRE_BoomerAMGSetNumSweeps`. Default: 2

linear_solvers.bamg_max_levels

See `HYPRE_BoomerAMGSetMaxLevels`. Default: 20

linear_solvers.bamg_strong_threshold

See `HYPRE_BoomerAMGSetStrongThreshold`. Default: 0.25

Time Integration Options**Time_Integrators**

A list of time-integration options used to advance the `realms` in time. Each list entry must contain a YAML mapping with the key indicating the type of time integrator. Currently only one option, `StandardTimeIntegrator` is available.

```
Time_Integrators:
- StandardTimeIntegrator:
  name: ti_1
  start_time: 0.0
  termination_step_count: 10
  time_step: 0.5
  time_stepping_type: fixed
  time_step_count: 0
  second_order_accuracy: yes

  realms:
  - fluids_realm
```

time_int.name

The lookup key for this time integration entry. This name must match the one provided in [Simulations](#) section.

time_int.termination_time

Nalu-Wind will stop the simulation once the `termination_time` has reached.

time_int.termination_step_count

Nalu-Wind will stop the simulation once the specified `termination_step_count` timesteps have been completed. If both `time_int.termination_time` and this parameter are provided then this parameter will prevail.

time_int.time_step

The time step (Δt) used for the simulation. If `time_int.time_stepping_type` is fixed this value does not change during the simulation.

time_int.start_time

The starting time step (default: 0.0) when starting a new simulation. Note that this has no effect on restart which is controlled by `restart.restart_time` in the `restart` section.

time_int.time_step_count

The starting timestep counter for a new simulation. See [restart](#) for restarting from a previous simulation.

time_int.second_order_accuracy

A boolean flag indicating whether second-order time integration scheme is activated. Default: no.

time_int.time_stepping_type

One of `fixed` or `adaptive` indicating whether a fixed time-stepping scheme or an adaptive timestepping scheme is used for simulations. See [time_step_control](#) for more information on max Courant number based adaptive time stepping.

time_int.realms

A list of `realms` names. The names entered here must match `name` used in the `realms` section. Names listed here not found in `realms` list will trigger an error, while realms not included in this list but present in `realms` will not be initialized and silently ignored. This can cause the code to abort if the user attempts to access the specific realm in the [transfers](#) section.

Physics Realm Options

As mentioned previously, `realms` is a YAML list data structure containing at least one [Physics Realm Options](#) entry that defines the computational domain (provided as an Exodus-II mesh), the set of physics equations that must be solved over this domain, along with the necessary initial and boundary conditions. Each list entry is a YAML dictionary mapping that is described in this section of the manual. The key subsections of a Realm entry in the input file are

Realm subsection	Purpose
<i>equation_systems</i>	Set of physics equations to be solved
<i>initial_conditions</i>	Initial conditions for the various fields
<i>boundary_conditions</i>	Boundary condition for the different fields
<i>material_properties</i>	Material properties (e.g., fluid density, viscosity etc.)
<i>solution_options</i>	Discretization and numerical stability
<i>mesh_transformation</i>	Mesh transformation
<i>mesh_motion</i>	Mesh motion
<i>output</i>	Solution output options (file, frequency, etc.)
<i>restart</i>	Optional: Restart options (restart time, checkpoint frequency etc.)
<i>time_step_control</i>	Optional: Parameters determining variable timestepping

In addition to the sections mentioned in the table, there are several additional sections that could be present depending on the specific simulation type and post-processing options requested by the user. A brief description of these optional sections are provided below:

Realm subsection	Purpose
<i>turbulence_averaging</i>	Generate statistics for the flow field
<i>post_processing</i>	Extract integrated data from the simulation
<i>solution_norm</i>	Compare the solution error to a reference solution
<i>data_probes</i>	Extract data using probes
<i>actuator</i>	Model turbine blades/tower using actuator lines
<i>abl_forcing</i>	Momentum source term to drive ABL flows to a desired velocity profile
<i>boundary_layer_statistics</i>	Compute boundary layer statistics

Common options

name

The name of the realm. The name provided here is used in the *Time_Integrators* section to determine the time-integration scheme used for this computational domain.

mesh

The name of the Exodus-II mesh file that defines the computational domain for this realm. Note that only the base name (i.e., without the `.NPROCS` `.IPROC` suffix) is provided even for simulations using previously decomposed mesh/restart files.

automatic_decomposition_type

Used only for parallel runs, this indicates how the a single mesh database must be decomposed amongst the MPI processes during initialization. This option should not be used if the mesh has already been decomposed by an external utility. Possible values are:

Value	Description
rcb	recursive coordinate bisection
rib	recursive inertial bisection
linear	elements in order first n/p to proc 0, next to proc 1.
cyclic	elements handed out to id % proc_count

activate_aura

A boolean flag indicating whether an extra element is *ghosted* across the processor boundaries. The default value is `no`.

use_edges

A boolean flag indicating whether edge based discretization scheme is used instead of element based schemes. The default value is `no`.

polynomial_order

An integer value indicating the polynomial order used for higher-order mesh simulations. The default value is 1. When *polynomial_order* is greater than 1, the Realm has the capability to promote the mesh to higher-order during initialization.

solve_frequency

An integer value indicating how often this realm is solved during time integration. The default value is 1.

support_inconsistent_multi_state_restart

A boolean flag indicating whether restarts are allowed from files where the necessary field states are missing. A typical situation is when the simulation is restarted using second-order time integration but the restart file was created using first-order time integration scheme.

activate_memory_diagnostic

A boolean flag indicating whether memory diagnostics are activated during simulation. Default value is `no`.

rebalance_mesh

A boolean flag indicating whether to rebalance mesh using `stk_balance`. The default value is `no`. If this parameter is activated, it requires that `stk_rebalance_method` is also set to specify the decomposition method to be used for rebalance, e.g., RIB, RCB, etc.

balance_nodes

A boolean flag indicating whether node balancing is performed during simulations. See also *balance_node_iterations* and *balance_nodes_target*.

balance_node_iterations

The frequency at which node rebalancing is performed. Default value is 5.

balance_node_target

The target balance ratio. Default value is 1.0.

Equation Systems

equation_systems

equation_systems subsection defines the physics equation sets that are solved for this realm and the linear solvers used to solve the different linear systems.

Note: The variable in the *equation_systems* subsection are prefixed with `equation_systems.name` but only the variable name after the period should appear in the input file.

equation_systems.name

A string indicating the name used in log messages etc.

equation_systems.max_iterations

The maximum number of non-linear iterations performed during a timestep that couples the different equation systems.

equation_systems.solver_system_specification

A mapping containing `field_name: linear_solver_name` that determines the linear solver used for solving the linear system. Example:

```

solver_system_specification:
  pressure: solve_continuity
  enthalpy: solve_scalar
  velocity: solve_scalar

```

The above example indicates that the linear systems for the enthalpy and momentum (velocity) equations are solved by the linear solver corresponding to the tag `solve_scalar` in the `linear_systems` entry, whereas the continuity equation system (pressure Poisson solve) should be solved using the linear solver definition corresponding to the tag `solve_continuity`.

equation_systems.systems

A list of equation systems to be solved within this realm. Each entry is a YAML mapping with the key corresponding to a pre-defined equation system name that contains additional parameters governing the solution of this equation set. The predefined equation types are

Equation system	Description
LowMachEOM	Low-Mach Momentum and Continuity equations
Enthalpy	Energy equations
ShearStressTransport	$k - \omega$ SST equation set
TurbKineticEnergy	TKE equation system
MassFraction	Mass Fraction
MixtureFraction	Mixture Fraction
MeshDisplacement	Arbitrary Mesh Displacement

An example of the equation system definition for ABL precursor simulations is shown below:

```

# Equation systems example for ABL precursor simulations
systems:
  - LowMachEOM:
      name: myLowMach
      max_iterations: 1
      convergence_tolerance: 1.0e-5
  - TurbKineticEnergy:
      name: myTke
      max_iterations: 1
      convergence_tolerance: 1.0e-2
  - Enthalpy:
      name: myEnth
      max_iterations: 1
      convergence_tolerance: 1.0e-2

```

Initial conditions

initial_conditions

The `initial_conditions` sub-sections defines the conditions used to initialize the computational fields if they are not provided via the mesh file. Two types of field initializations are currently possible:

- `constant` - Initialize the field with a constant value throughout the domain;
- `user_function` - Initialize the field with a pre-defined user function.

The snippet below shows an example of both options available to initialize the various computational fields used for ABL simulations. In this example, the pressure and turbulent kinetic energy fields are initialized using a constant value, whereas the velocity field is initialized by the user function `boundary_layer_perturbation` that imposes sinusoidal fluctuations over a velocity field to trip the flow.

```
initial_conditions:
- constant: ic_1
  target_name: [fluid_part]
  value:
    pressure: 0.0
    turbulent_ke: 0.1

- user_function: ic_2
  target_name: [fluid_part]
  user_function_name:
    velocity: boundary_layer_perturbation
  user_function_parameters:
    velocity: [1.0, 0.0075398, 0.0075398, 50.0, 8.0]
```

initial_conditions.constant

This input parameter serves two purposes: 1. it indicates the *type* (constant), and 2. provides the custom *name* for this condition. In addition to the `initial_conditions.target_name` this section requires another entry *value* that contains the mapping of (field_name, value) as shown in the above example.

initial_conditions.user_function

Indicates that this block of YAML input must be parsed as input for a user defined function.

initial_conditions.target_name

A list of element blocks (*parts*) where this initial condition must be applied. Using the alias `all_blocks` is equivalent to listing all element blocks in the mesh.

Boundary Conditions

boundary_conditions

This subsection of the physics Realm contains a list of boundary conditions that must be used during the simulation. Each entry of this list is a YAML mapping entry with the key of the form `<type>_boundary_condition` where the available types are:

- inflow
- open – Outflow BC
- wall
- symmetry
- periodic
- non_conformal – e.g., BC across sliding mesh interfaces
- overset – overset mesh assembly description

All BC types require `bc.target_name` that contains a list of side sets where the specified BC is to be applied. Additional information necessary for certain BC types are provided by a sub-dictionary with the key `<type>_user_data:` that contains the parameters necessary to initialize a specific BC type.

bc.target_name

A list of side set part names where the given BC type must be applied. If a single string value is provided, it is converted to a list internally during input file processing phase.

Inflow Boundary Condition

```
- inflow_boundary_condition: bc_inflow
  target_name: inlet
  inflow_user_data:
    velocity: [0.0,0.0,1.0]
```

Open Boundary Condition

```
- open_boundary_condition: bc_open
  target_name: outlet
  open_user_data:
    velocity: [0,0,0]
    pressure: 0.0
    entrainment_method: {computed, specified}
    total_pressure: {yes, no}
```

Wall Boundary Condition

`bc.wall_user_data`

This subsection contains specifications as to whether wall models are used, or how to treat the velocity at the wall when there is mesh motion.

The following input file snippet shows an example of using an ABL wall function at the terrain during ABL simulations. See [ABL Wall Function](#) for more details on the actual implementation.

```
# Wall boundary condition example for ABL terrain modeling
- wall_boundary_condition: bc_terrain
  target_name: terrain
  wall_user_data:
    velocity: [0,0,0]
    use_abl_wall_function: yes
    heat_flux: 0.0
    roughness_height: 0.2
    gravity_vector_component: 3
    reference_temperature: 300.0
```

The entry `gravity_vector_component` is an integer that specifies the component of the gravity vector, defined in `solution_options.gravity`, that should be used in the definition of the Monin-Obukhov length scale calculation. The entry `reference_temperature` is the reference temperature used in calculation of the Monin-Obukhov length scale.

When there is mesh motion involved the wall boundary velocity takes the value of the `mesh_velocity` along the part represented by `bc.target_name`. In such a scenario all information under `bc.wall_user_data` is rendered unused.

Example of wall boundary with a custom user function for temperature at the wall

```
- wall_boundary_condition: bc_6
  target_name: surface_6
  wall_user_data:
    user_function_name:
      temperature: steady_2d_thermal
```

Symmetry Boundary Condition

Requires no additional input other than `bc.target_name`.

```
- symmetry_boundary_condition: bc_top
  target_name: top
  symmetry_user_data:
```

Periodic Boundary Condition

Unlike the other BCs described so far, the parameter `bc.target_name` behaves differently for the periodic BC. This parameter must be a list containing exactly two entries: the boundary face pair where periodicity is enforced. The nodes on these planes must coincide after translation in the direction of periodicity. This BC also requires a `periodic_user_data` section that specifies the search tolerance for locating node pairs.

periodic_user_data

```
- periodic_boundary_condition: bc_east_west
  target_name: [east, west]
  periodic_user_data:
    search_tolerance: 0.0001
```

Non-Conformal Boundary

Like the periodic BC, the parameter `bc.target_name` must be a list with exactly two entries that specify the boundary plane pair forming the non-conformal boundary.

```
- non_conformal_boundary_condition: bc_left
  target_name: [surface_77, surface_7]
  non_conformal_user_data:
    expand_box_percentage: 10.0
```

Material Properties

material_properties

The section provides the properties required for various physical quantities during the simulation. A sample section used for simulating ABL flows is shown below

```
material_properties:
  target_name: [fluid_part]

  constant_specification:
    universal_gas_constant: 8314.4621
    reference_pressure: 101325.0

  reference_quantities:
    - species_name: Air
      mw: 29.0
      mass_fraction: 1.0

  specifications:
```

(continues on next page)

(continued from previous page)

```

- name: density
  type: constant
  value: 1.178037722969475
- name: viscosity
  type: constant
  value: 1.6e-5
- name: specific_heat
  type: constant
  value: 1000.0

```

material_properties.target_name

A list of element blocks (*parts*) where the material properties are applied. This list should ideally include all the parts that are referenced by `initial_conditions.target_name`. Using the alias `all_blocks` is equivalent to listing all element blocks in the mesh.

material_properties.constant_specification

Values for several constants used during the simulation. Currently the following properties are defined:

Name	Description
universal_gas_constant	Ideal gas constant R
reference_temperature	Reference temperature for simulations
reference_pressure	Reference pressure for simulations

material_properties.reference_quantities

Provides material properties for the different species involved in the simulation.

Name	Description
species_name	Name used to lookup properties
mw	Molecular weight
mass_fraction	Mass fraction
primary_mass_fraction	
secondary_mass_fraction	
stoichiometry	

material_properties.specifications

A list of material properties with the following parameters

material_properties.specifications.name

The name used for lookup, e.g., density, viscosity, etc.

material_properties.specifications.type

The type can be one of the following

Type	Description
constant	Constant value property
polynomial	Property determined by a polynomial function
ideal_gas_t	Function of T_{ref} , P_{ref} , molecular weight
ideal_gas_t_p	Function of T_{ref} , pressure, molecular weight
ideal_gas_yk	
hdf5table	Lookup from an HDF5 table
mixture_fraction	Property determined by the mixture fraction
geometric	
generic	

Examples

1. Specification for density as a function of temperature

```
specifications:
- name: density
  type: ideal_gas_t
```

2. Specification of viscosity as a function of temperature

```
- name: viscosity
  type: polynomial
  coefficient_declaration:
  - species_name: Air
    coefficients: [1.7894e-5, 273.11, 110.56]
```

The `species_name` must correspond to the entry in *reference quantities* to lookup molecular weight information.

3. Specification via hdf5table

```
material_properties:
  table_file_name: SLFM_CGauss_C2H4_ZMean_ZScaledVarianceMean_logChiMean.h5

  specifications:
  - name: density
    type: hdf5table
    independent_variable_set: [mixture_fraction, scalar_variance, scalar_
↪dissipation]
    table_name_for_property: density
    table_name_for_independent_variable_set: [ZMean, ZScaledVarianceMean,
↪ChiMean]
    aux_variables: temperature
    table_name_for_aux_variables: temperature

  - name: viscosity
    type: hdf5table
    independent_variable_set: [mixture_fraction, scalar_variance, scalar_
↪dissipation]
    table_name_for_property: mu
    table_name_for_independent_variable_set: [ZMean, ZScaledVarianceMean,
↪ChiMean]
```

4. Specification via mixture_fraction

```
material_properties:
  target_name: block_1

  specifications:
  - name: density
    type: mixture_fraction
    primary_value: 0.163e-3
    secondary_value: 1.18e-3
  - name: viscosity
    type: mixture_fraction
    primary_value: 1.967e-4
    secondary_value: 1.85e-4
```


Solution Options

Note: The documentation for this section is incomplete.

solution_options

This section defines the discretization and numerical stability approaches, as well as turbulence models.

solution_options.name

Name of solution options group.

solution_options.turbulence_model

Turbulence model used in simulation.

solution_options.options

This subsection defines additional options for the solution options.

For example, one could modify turbulence model constants:

```
- turbulence_model_constants:
  SDRWallFactor: 0.625
```

One could also define source terms, such as a momentum forcing in a box of the domain:

```
- source_terms:
  momentum: body_force_box

- source_term_parameters:
  momentum: [0.011, 0.0, 0.0]
  momentum_box: [-1.0, 1.00001, 0.0, 10.0, 4.0, 5.0]
```

Mesh Transformation

mesh_transformation

This subsection of the realm describes a one time stationary motion undergone by the entire mesh with entries under *mesh_transformation* describing the motions applied to different parts in a.

Example:

```
mesh_transformation:
- name: scale_background
  mesh_parts: [ Unspecified-3-HEX ]
  motion:
    - type: scaling
      factor: [1.2, 1.0, 1.2]
      origin: [5.0, 0.05, 0.0]

- name: scale_near_body
  mesh_parts: [ Unspecified-2-HEX ]
  motion:
    - type: scaling
      factor: [1.2, 1.0, 1.2]
      origin: [0.0, 0.05, 0.0]
```

mesh_transformation.name

Name of motion group.

mesh_transformation.mesh_parts

Mesh parts associated with respective motion group. The user may use `all_blocks` to apply the transformation to the entire mesh.

mesh_transformation.motion

Type of motion. Every group is free to undergo one or multiple motions simultaneously.

Mesh Motion

mesh_motion

This subsection of the of the realm describes the time-dependent rigid body motion undergone by the entire mesh for as described by entries under `mesh_motion`.

Example:

```
mesh_motion:
- name: trans_rot_near_body
  mesh_parts: [ Unspecified-2-HEX ]
  motion:
    - type: rotation
      omega: 12.0
      axis: [0.0, 1.0, 0.0]
      origin: [0.0, 0.05, 0.0]

    - type: translation
      start_time: 100.0
      end_time: 200.0
      velocity: [0.05, 0.0, 0.0]
```

mesh_motion.name

Name of motion group.

mesh_motion.mesh_parts

Mesh parts associated with respective motion group. The user may use `all_blocks` to apply the motion to the entire mesh.

mesh_motion.motion

Type of motion the current group undergoes. Every frame is free to undergo one or multiple motions simultaneously.

Output Options

output

Specifies the frequency of output, the output database name, etc.

Example:

```
output:
  output_data_base_name: out/ABL.neutral.e
  output_frequency: 100
  output_node_set: no
  output_variables:
    - velocity
    - pressure
    - temperature
```

output.output_data_base_name

The name of the output Exodus-II database. Can specify a directory relative to the run directory, e.g., `out/nalu_results.e`. The directory will be created automatically if one doesn't exist. Default: `output.e`

output.output_frequency

Nalu-Wind will write the output file every `output_frequency` timesteps. Note that currently there is no option to output results at a specified simulation time. Default: 1.

output.output_start

Nalu-Wind will start writing output past the `output_start` timestep. Default: 0.

output.output_forced_wall_time

Force output at a specified *wall-clock time* in seconds.

output.output_node_set

Boolean flag indicating whether nodesets, if present, should be output to the output file along with element blocks.

output.compression_level

Integer value indicating the compression level used. Default: 0.

output.output_variables

A list of field names to be output to the database. The field variables can be node or element based quantities.

Restart Options

restart

This section manages the restart for this realm object.

restart.restart_data_base_name

The filename for restart. Like `output`, the filename can contain a directory and it will be created if not already present.

restart.restart_time

If this variable is present, it indicates that the current run will restart from a previous simulation. This requires that the `mesh` be a restart file with all the fields necessary for the equation sets defined in the `equation_systems.systems`. Nalu-Wind will restart from the closest time available in the `mesh` to `restart_time`. The timesteps available in a restart file can be examined by looking at the `time_whole` variable using the `ncdump` utility.

Note: The restart database used for restarting a simulation is the `mesh` parameter. The `restart_data_base_name` parameter is used exclusively for outputs.

restart.restart_frequency

The frequency at which restart files are written to the disk. Default: 500 timesteps.

restart.restart_start

Nalu-Wind will write a restart file after `restart_start` timesteps have elapsed.

restart.restart_forced_wall_time

Force writing of restart file after specified *wall-clock time* in seconds.

restart.restart_node_set

A boolean flag indicating whether nodesets are output to the restart database.

restart.max_data_base_step_size

Default: 100,000.

restart.compression_level

Compression level. Default: 0.

Time-step Control Options

time_step_control

This optional section specifies the adaptive time stepping parameters used if *time_int.time_stepping_type* is set to adaptive.

```
time_step_control:
  target_courant: 2.0
  time_step_change_factor: 1.2
```

dtctrl.target_courant

Maximum Courant number allowed during the simulation. Default: 1.0

dtctrl.time_step_change_factor

Maximum allowable increase in dt over a given timestep.

Turbine specific input options

Actuator Turbine Model

actuator

actuator subsection defines the inputs for actuator line simulations. A sample section is shown below for running actuator line simulations coupled to OpenFAST with two turbines.

```
actuator:
  type: ActLineFAST
  search_method: stk_kdtree
  search_target_part: Unspecified-2-HEX

  n_turbines_glob: 2
  dry_run: False
  debug: False
  t_start: 0.0
  simStart: init # init/trueRestart/restartDriverInitFAST
  t_max: 5.0
  n_every_checkpoint: 100

  Turbine0:
    procNo: 0
    num_force_pts_blade: 50
    num_force_pts_tower: 20
    nacelle_cd: 1.0
    nacelle_area: 1.0
    air_density: 1.225
    epsilon: [ 5.0, 5.0, 5.0 ]
    turbine_base_pos: [ 0.0, 0.0, -90.0 ]
    turbine_hub_pos: [ 0.0, 0.0, 0.0 ]
    restart_filename: ""
    FAST_input_filename: "Test01.fst"
    turb_id: 1
    turbine_name: machine_zero
```

(continues on next page)

(continued from previous page)

```

Turbine1:
  procNo: 0
  num_force_pts_blade: 50
  num_force_pts_tower: 20
  nacelle_cd: 1.0
  nacelle_area: 1.0
  air_density: 1.225
  epsilon: [ 5.0, 5.0, 5.0 ]
  turbine_base_pos: [ 250.0, 0.0, -90.0 ]
  turbine_hub_pos: [ 250.0, 0.0, 0.0 ]
  restart_filename: ""
  FAST_input_filename: "Test02.fst"
  turb_id: 2
  turbine_name: machine_one

```

actuator.type

Type of actuator source. Options are ActLineFAST and ActDiskFAST. ActLineFAST is for actuator lines, and ActDiskFAST is for actuator disks. The actuator disk uses a stationary actuator line model to compute forces at the blade locations and then the average force of the blades is spread azimuthally between the blades sampling points.

actuator.search_method

String specifying the type of search method used to identify the nodes within the search radius of the actuator points. The only valid option is `stk_kdtree`. The `boost_rtree` option has been deprecated by the STK search library.

search_target_part

String or an array of strings specifying the parts of the mesh to be searched to identify the nodes near the actuator points.

actuator.n_turbines_glob

Total number of turbines in the simulation. The input file must contain a number of turbine specific sections (*Turbine0*, *Turbine1*, ..., *Turbine(n-1)*) that is consistent with *nTurbinesGlob*.

actuator.debug

Enable debug outputs if set to true

actuator.dry_run

The simulation will not run if `dryRun` is set to true. However, the simulation will read the input files, allocate turbines to processors and prepare to run the individual turbine instances. This flag is useful to test the setup of the simulation before running it.

actuator.simStart

Flag indicating whether the simulation starts from scratch or restart. `simStart` takes on one of three values:

- `init` - Use this option when starting a simulation from $t=0s$.
- `trueRestart` - While OpenFAST allows for restart of a turbine simulation, external components like the Bladed style controller may not. Use this option when all components of the simulation are known to restart.
- `restartDriverInitFAST` - When the `restartDriverInitFAST` option is selected, the individual turbine models start from $t=0s$ and run up to the specified restart time using the inflow data stored at the actuator nodes from a hdf5 file. The C++ API stores the inflow data at the actuator nodes in a hdf5 file at every OpenFAST time step and then reads it back when using this restart option. This restart option is especially useful when the glue code is a CFD solver.

actuator.t_start

Start time of the simulation

actuator.t_end

End time of the simulation. $t_{end} \leq t_{max}$

actuator.t_max

Max time of the simulation

Note: t_{max} can only be set when OpenFAST is running from $t=0s$ and `simStart` is `init`. t_{max} can not be changed on a restart. OpenFAST will not be able to run beyond t_{max} . Choose t_{max} to be large enough to accomodate any possible future extensions of runs. One can change t_{start} and t_{end} to start and stop the simulation any number of times as long as $t_{end} \leq t_{max}$.

actuator.dt_fast

Time step for OpenFAST. All turbines should have the same time step.

actuator.n_every_checkpoint

Restart files will be written every so many time steps

Turbine specific input options

actuator.turbine_base_pos

The position of the turbine base for actuator-line/disk simulations

actuator.num_force_pts_blade

The number of actuator points along each blade for actuator-line/disk simulations

actuator.num_force_pts_tower

The number of actuator points along the tower for actuator-line/disk simulations.

actuator.nacelle_cd

The drag coefficient for the nacelle. If this is set to zero, or not defined, the code will not implement the nacelle model.

actuator.nacelle_area

The reference area for the nacelle. This is only used if the nacelle model is used.

actuator.air_density

The air density. This is only used to compute the nacelle force. It should match the density being used in both Nalu and OpenFAST.

actuator.epsilon

The spreading width ϵ in the Gaussian spreading function in the *[chordwise, thickness, spanwise]* coordinate system to spread the forces from the actuator point to the nodes. In the case of the actuator disk, only the first value in the chordwise direction is used for the uniform isotropic Gaussian.

actuator.epsilon_chord

This is the ratio ϵ/c in every direction *[chordwise, thickness, spanwise]*. If this option is specified, the code will choose a value of ϵ at every location that is $c * \epsilon/c$. To avoid numerical instabilities, the code will choose the maximum value between $c * \epsilon/c$ and the value of `actuator.epsilon_min` specified.

actuator.epsilon_min

This is the minimum value of ϵ in the Gaussian spreading function in the *[chordwise, thickness, spanwise]* coordinate system to spread the forces from the actuator point to the nodes. This option is required if the option `actuator.epsilon_chord` is specified.

actuator.epsilon_tower

The spreading width ϵ in the Gaussian spreading function in the inertial *[x, y, z]* reference frame. If this value is not specified, then `actuator.epsilon` or `actuator.epsilon_min` will be used.

actuator.restart_filename

The checkpoint file for this turbine when restarting a simulation

actuator.FAST_input_filename

The FAST input file for this turbine

actuator.turb_id

A unique turbine id for each turbine

actuator.num_swept_pts

This is an optional parameter specifically for actuator disks. This parameter determines the number of points that are placed azimuthally between the actuator lines and spread the forcing over the disk's area. When `num_swept_pts` is included the number of azimuthal points between the lines is forced to this value at all radial locations. If `num_swept_pts` is omitted then the azimuthal sampling is computed automatically with different sampling at each radial location such that the average distance between points matches the radial spacing.

Turbulence averaging

turbulence_averaging

`turbulence_averaging` subsection defines the turbulence post-processing quantities and averaging procedures. A sample section is shown below

```
turbulence_averaging:
  forced_reset: no
  time_filter_interval: 100000.0

  averaging_type: nalu_classic/moving_exponential

  specifications:
    - name: turbulence_postprocessing
      target_name: interior
      reynolds_averaged_variables:
        - velocity

      favre_averaged_variables:
        - velocity
        - resolved_turbulent_ke

      compute_tke: yes
      compute_reynolds_stress: yes
      compute_resolved_stress: yes
      compute_temperature_resolved_flux: yes
      compute_sfs_stress: yes
      compute_temperature_sfs_flux: yes
      compute_q_criterion: yes
      compute_vorticity: yes
      compute_lambda_ci: yes
```

Note: The variable in the `turbulence_averaging` subsection are prefixed with `turbulence_averaging.` name but only the variable name after the period should appear in the input file.

turbulence_averaging.forced_reset

A boolean flag indicating whether the averaging of all quantities in the turbulence averaging section is reset. If this flag is true, the running average is set to zero.

turbulence_averaging.averaging_type

This parameter sets the choice of the running average type. Possible values are:

nalu_classic “Sawtooth” average. The running average is set to zero each time the time filter width is reached and a new average is calculated for the next time interval.

moving_exponential “Moving window” average where the window size is set to to the time filter width. The contribution of any quantity before the moving window towards the average value reduces exponentially with every time step.

turbulence_averaging.time_filter_interval

Number indicating the time filter size over which to calculate the running average. This quantity is used in different ways for each filter discussed above.

turbulence_averaging.specifications

A list of turbulence postprocessing properties with the following parameters

turbulence_averaging.specifications.name

The name used for lookup and logging.

turbulence_averaging.specifications.target_name

A list of element blocks (parts) where the turbulence averaging is applied.

turbulence_averaging.specifications.reynolds_average_variables

A list of field names to be averaged.

turbulence_averaging.specifications.favre_average_variables

A list of field names to be Favre averaged.

turbulence_averaging.specifications.compute_tke

A boolean flag indicating whether the turbulent kinetic energy is computed. The default value is `no`.

turbulence_averaging.specifications.compute_reynolds_stress

A boolean flag indicating whether the reynolds stress is computed. The default value is `no`.

turbulence_averaging.specifications.compute_resolved_stress

A boolean flag indicating whether the average resolved stress is computed as $\langle \bar{\rho} \tilde{u}_i \tilde{u}_j \rangle$. The default value is `no`. When this option is turned on, the Favre average of the resolved velocity, $\langle \bar{\rho} \tilde{u}_j \rangle$, is computed as well.

turbulence_averaging.specifications.compute_temperature_resolved_flux

A boolean flag indicating whether the average resolved temperature flux is computed as $\langle \bar{\rho} \tilde{u}_i \tilde{\theta} \rangle$. The default value is `no`. When this option is turned on, the Favre average of the resolved temperature, $\langle \bar{\rho} \tilde{\theta} \rangle$, is computed as well.

turbulence_averaging.specifications.compute_sfs_stress

A boolean flag indicating whether the average sub-filter scale stress is computed. The default value is `no`. The sub-filter scale stress model is assumed to be of an eddy viscosity type and the turbulent viscosity computed by the turbulence model is used. The sub-filter scale kinetic energy is used to determine the isotropic component of the sub-filter stress. As described in the section *Conservation of Momentum*, the Yoshizawa model is used to compute the sub-filter kinetic energy when it is not transported.

turbulence_averaging.specifications.compute_temperature_sfs_flux

A boolean flag indicating whether the average sub-filter scale flux of temperature is computed. The default value is `no`. The sub-filter scale stress model is assumed to be of an eddy diffusivity type and the turbulent diffusivity computed by the turbulence model is used along with a constant turbulent Prandtl number obtained from the Realm.

turbulence_averaging.specifications.compute_favre_stress

A boolean flag indicating whether the Favre stress is computed. The default value is `no`.

turbulence_averaging.specifications.compute_favre_tke

A boolean flag indicating whether the Favre stress is computed. The default value is `no`.

turbulence_averaging.specifications.compute_q_criterion

A boolean flag indicating whether the q-criterion is computed. The default value is `no`.

turbulence_averaging.specifications.compute_vorticity

A boolean flag indicating whether the vorticity is computed. The default value is `no`.

turbulence_averaging.specifications.compute_lambda_ci

A boolean flag indicating whether the Lambda2 vorticity criterion is computed. The default value is `no`.

Data probes

data_probes

`data_probes` subsection defines the data probes. A sample section is shown below

```
data_probes:
  output_frequency: 100
  output_format: text
  search_method: stk_octree
  search_tolerance: 1.0e-3
  search_expansion_factor: 2.0

  gzip_level: 0
  write_coords: true

  specifications:
    - name: probe_bottomwall
      from_target_part: bottomwall

      line_of_site_specifications:
        - name: probe_bottomwall
          number_of_points: 100
          tip_coordinates: [-6.39, 0.0, 0.0]
          tail_coordinates: [4.0, 0.0, 0.0]

      output_variables:
        - field_name: tau_wall
          field_size: 1
        - field_name: pressure

  specifications:
    - name: probe_profile
      from_target_part: interior

      line_of_site_specifications:
        - name: probe_profile
          number_of_points: 100
          tip_coordinates: [0, 0.0, 0.0]
          tail_coordinates: [0.0, 0.0, 1.0]

      plane_specifications:
        - name: sample_plane
          corner_coordinates: [0.0, 0.0, 0.0]
          edge1_vector: [1.0, 0, 0]
          edge2_vector: [0, 2.0, 0]
          edge1_numPoints: 11
          edge2_numPoints: 21
          offset_vector: [0, 0, 1]
```

(continues on next page)

(continued from previous page)

```
offset_spacings: [0, 2]
only_output_field: velocity

output_variables:
- field_name: velocity
  field_size: 3
- field_name: reynolds_stress
  field_size: 6
```

Note: The variable in the `data_probes` subsection are prefixed with `data_probes.name` but only the variable name after the period should appear in the input file.

data_probes.output_frequency

Integer specifying the frequency of output.

data_probes.output_format

String specifying the output format for the data probes. Currently available options are `text` or `exodus`. If not specified, the default is `text`. Multiple output formats can be specified like the following:

```
output_format:
- text
- exodus
```

data_probes.search_method

String specifying the search method for finding nodes to transfer field quantities to the data probe lineout.

data_probes.search_tolerance

Number specifying the search tolerance for locating nodes.

data_probes.search_expansion_factor

Number specifying the factor to use when expanding the node search.

data_probes.gzip_level

Optional input, applies to sample planes only. Integer specifying amount of compression to apply to sample plane output. The default `gzip_level=0`, means no compression. To apply compression, use `gzip_level` from 1 to 9, with 9 indicating maximum compression (and slowest speed). Generally `gzip_level=1` or `gzip_level=2` is sufficient.

data_probes.write_coords

Optional input, applies to sample planes only. Boolean specifying whether the sample plane x,y,z coordinates and indices are to be included with every sample plane output. The default is `write_coords=true`. For `write_coords=false`, a separate coordinate file will be written at the beginning of the output sequence if it does not already exist.

data_probes.time_performance

Optional input, applies to sample planes only. Boolean specifying whether to display timing information when writing sample planes.

data_probes.specifications

A list of data probe properties with the following parameters

data_probes.specifications.name

The name used for lookup and logging.

data_probes.specifications.from_target_part

A list of element blocks (parts) where to do the data probing.

data_probes.specifications.line_of_site_specifications

A list specifications defining the lineout

Parameter	Description
name	File name (without extension) for the data probe
number_of_points	Number of points along the lineout
tip_coordinates	List containing the coordinates for the start of the lineout
tail_coordinates	List containing the coordinates for the end of the lineout

data_probes.specifications.plane_specifications

A list specifications defining the sampling plane

Parameter	Description
name	File name (without extension) for the sampling plane
corner_coordinates	List containing the coordinates for the corner of the plane
edge1_vector	List containing the vector defining the first edge of the plane (with origin at corner)
edge2_vector	List containing the vector defining the second edge of the plane (with origin at corner)
edge1_numPoints	Number of points along edge 1
edge2_numPoints	Number of points along edge 2
offset_vector	[Optional] List containing the vector defining the offset direction for additional planes
offset_spacings	[Optional] List containing how far each plane is to be offset in the offset_vector direction
only_output_field	[Optional] Only include the output of this variable in the sample plane output.

data_probes.specifications.output_variables

A list of field names (and field size) to be probed.

data_probes.lidar_specifications

Allows line_of_site sampling along trajectories tracing the rosette pattern of a spinner LIDAR.

data_probes.lidar_specifications.from_target_part

The mesh part containing the spinner LIDAR center coordinates.

data_probes.lidar_specifications.scan_time

The time for a scan by the simulated spinner LIDAR.

data_probes.lidar_specifications.number_of_samples

The number of lines generated by the spinner LIDAR sampling. For the text output, this will generate a separate file for each line.

data_probes.lidar_specifications.points_along_line

The number samples along each lines. This should be chosen based on the spatial resolution of the underlying mesh, the LIDAR, measurements and the *beam_length* parameter.

data_probes.lidar_specifications.center

The location of the spinner LIDAR aperture.

data_probes.lidar_specifications.beam_length

The maximum length over which to sample the velocity on a particular line. The spatial resolution of the sampling is computed from this and the *number_of_samples* parameter.

data_probes.lidar_specifications.axis

The orientation vector for the LIDAR measurements.

dataprob.es.lidar_specifications.misc

The user may also set a number of parameters corresponding to the hardware configuration of the spinner LIDAR.

Parameter	Description
inner_prism_theta	Default 90 degrees. The starting angle of the inner prism
inner_prism_rotation_rate	Default 3.5 degrees per second. Rotation rate of the inner prism
inner_prism_azimuth	Default 15.2 degrees. azimuthal angle of the inner prism
outer_prism_theta	Default 90 degrees. The starting angle of the outer prism
outer_prism_rotation_rate	Default 6.5 degrees per second. Rotation rate of the outer prism
outer_prism_azimuth	Default 15.2 degrees. azimuthal angle of the outer prism
ground_direction	Default [0,0,1]. Orthogonal orientation vector for the LIDAR

Post-processing**post_processing**

`post_processing` subsection defines the different post-processing options. A sample section is shown below

```
post_processing:
- type: surface
  physics: surface_force_and_moment
  output_file_name: results/wallHump.dat
  frequency: 100
  parameters: [0,0]
  target_name: bottomwall
```

Note: The variable in the `post_processing` subsection are prefixed with `post_processing.name` but only the variable name after the period should appear in the input file.

post_processing.type

Type of post-processing. Possible values are:

Value	Description
surface	Post-processing of surface quantities

post_processing.physics

Physics to be post-processing. Possible values are:

Value	Description
surface_force_and_moment	Calculate surface forces and moments
surface_force_and_moment_wall_function	Calculate surface forces and moments when using a wall function

post_processing.output_file_name

String specifying the output file name.

post_processing.frequency

Integer specifying the frequency of output.

post_processing.parameters

Parameters for the physics function. For the `surface_force_and_moment` type functions, this is a list specifying the centroid coordinates used in the moment calculation.

post_processing.target_name

A list of element blocks (parts) where to do the post-processing

ABL Forcing**abl_forcing**

`abl_forcing` allows the user to specify desired velocities and temperatures at different heights. These velocities and temperatures are enforced through the use of source in the momentum and enthalpy equations. The `abl_forcing` option needs to be specified in the momentum and/or enthalpy source blocks:

```
- source_terms:
  momentum: abl_forcing
  enthalpy: abl_forcing
```

This option allows the code to implement source terms in the momentum and/or enthalpy equations. A sample section is shown below

```
abl_forcing:
  search_method: stk_kdtree
  search_tolerance: 0.0001
  search_expansion_factor: 1.5
  output_frequency: 1

  from_target_part: [fluid_part]

  momentum:
    type: computed
    relaxation_factor: 1.0
    heights: [250.0, 500.0, 750.0]
    target_part_format: "zplane_%06.1f"

    # The velocities at each plane
    # Each list include a time and the velocities for each plane
    # Notice that the total number of elements in each list will be
    # number of planes + 1
    velocity_x:
      - [0.0, 10.0, 5.0, 15.0]
      - [100000.0, 10.0, 5.0, 15.0]

    velocity_y:
      - [0.0, 0.0, 0.0, 0.0]
      - [100000.0, 0.0, 0.0, 0.0]

    velocity_z:
      - [0.0, 0.0, 0.0, 0.0]
      - [100000.0, 0.0, 0.0, 0.0]

  temperature:
    type: computed
    relaxation_factor: 1.0
    heights: [250.0, 500.0, 750.0]
    target_part_format: "zplane_%06.1f"
```

(continues on next page)

(continued from previous page)

```
temperature:
- [0.0, 300.0, 280.0, 310.0]
- [100000.0, 300.0, 280.0, 310.0]
```

Note: The variables in the `abl_forcing` subsection are prefixed with `abl_forcing.name` but only the variable name after the period should appear in the input file.

abl_forcing.search_method

This specifies the search method algorithm within the stk framework. The default option `stk_kdtree` is recommended.

abl_forcing.search_tolerance

This is the tolerance specified for the `search_method` algorithm. A default value of 0.0001 is recommended.

abl_forcing.search_expansion_factor

This option is related to the stk search algorithm. A value of 1.5 is recommended.

abl_forcing.output_frequency

This is the frequency at which the source term is written to the output value. A value of 1 means the source term will be written to the output file every time-step.

Note: There are now two options in the following inputs. The can be `momentum` and/or `temperature`.

abl_forcing.momentum.computed

This option allows the user to choose if a momentum source is computed from a desired velocity (`computed`) or if a user defined source term is directly applied into the momentum equation (`user_defined`).

abl_forcing.momentum.relaxation_factor

This is a relaxation factor which can be used to under/over-relax the momentum source term. The default value is 1.

abl_forcing.momentum.heights

This is a list containing the planes at which the forcing should be implemented. Each input is the height for that plane. This is the naming convention in the mesh file.

abl_forcing.momentum.target_part_format

This is the format in which the planes are saved in the mesh file.

abl_forcing.momentum.velocity_x

A set of lists containing the time in the first element, followed by the desired velocity at each plane in the *x* direction.

abl_forcing.momentum.velocity_y

A set of lists containing the time in the first element, followed by the desired velocity at each plane in the *y* direction.

abl_forcing.momentum.velocity_z

A set of lists containing the time in the first element, followed by the desired velocity at each plane in the *z* direction.

Note: The temperature has the same inputs as the momentum source (`abl_forcing.temperature.type`, `abl_forcing.temperature.relaxation_factor`, `abl_forcing.temperature.heights`,

and `abl_forcing.temperature.target_part_format`) which take the same options.

`abl_forcing.temperature.temperature`

A set of lists containing the time in the first element, followed by the desired temperature at each plane.

Boundary Layer Statistics

`boundary_layer_statistics`

The `boundary_layer_statistics` subsection defines the statistics to be gathered from the ABL precursor calculation. This section computes the spatial averages of velocity and (optionally) temperature at all height levels available in the ABL mesh.

The outputs are a series of text files (`abl_*_stats.dat`) containing the averaged profiles and a netcdf file (e.g., `abl_statistics.nc`) containing the time history of the averaged quantities.

A sample section is shown below:

```
boundary_layer_statistics:
  target_name: [fluid_part]
  stats_output_file: abl_statistics.nc
  compute_temperature_statistics: yes
  output_frequency: 5000
  time_hist_output_frequency: 1
  height_multiplier: 1.0e6
```

The various parameters to `boundary_layer_statistics` are described below:

`boundary_layer_statistics.target_name`

A list of element blocks (*parts*) where the ABL statistics are to be computed.

`boundary_layer_statistics.time_filter_interval`

The length of time, in seconds, over which to average the statistics given in the `abl_*_stats.dat` files. [Optional, default value: 3600.0]

`boundary_layer_statistics.compute_temperature_statistics`

A yes or no value which indicates whether to include the averaged temperature statistics. [Optional, default value: yes]

`boundary_layer_statistics.output_frequency`

The frequency to output statistics in the `abl_*_stats.dat` text files. [Optional, default value: 10]

`boundary_layer_statistics.time_hist_output_frequency`

The frequency, in iterations, of the time history statistics included in the netcdf statistics file. [Optional, default value: 10]

`boundary_layer_statistics.stats_output_file`

The name of the netcdf statistics file which includes the time history and averages. [Optional, default value: `abl_statistics.nc`]

`boundary_layer_statistics.process_utau_statistics`

A yes or no value to indicate whether the utau statistics are to be included in the computations. [Optional, default value: yes]

`boundary_layer_statistics.wall_normal_direction`

Spatial index to indicate the wall normal direction in the domain. The directions are given by `x="1"`, `y="2"`, `z="3"`. [Optional, default value: 3]

`boundary_layer_statistics.minimum_height`

Minimum height to account for negative values in the wall normal direction. [Optional, default value: 0.0]

boundary_layer_statistics.height_multiplier

For the purposes of determining the unique heights for the ABL statistics, wall normal distances are multiplied by `height_multiplier` then converted into integers for binning. Larger values of `height_multiplier` allow a higher precision to be used in determining the unique heights and better behavior in some meshes. [Optional, default value: 1.0e6]

1.2.4 Transfers

transfers

Transfers section describes the search and mapping operations to be performed between participating realms within a simulation.

1.2.5 Simulations

simulations

This is the top-level section that orchestrates the entire execution of Nalu-Wind.

1.2.6 Lessons Learned from Meshing the McAlister Case

Author Chris Bruner, Dept. 01515, Sandia National Laboratories

Introduction

The series of wind-tunnel tests described by McAlister & Takahashi [McAl1991] have become something of a canonical test case in the rotorcraft community. This is because the tests are well-documented and investigate both tip and aspect ratio effects, and because the symmetric wing section used is fairly representative of those typically found on rotorcraft.

This case also serves as a reasonably good test case for wind energy applications as there are measurements of the trailing tip vortex far downstream, up to 13 chords. This is important to understand the grid requirements of our unstructured approach to modeling a full-scale blade-resolved rotor and tower system.

Meshes

The meshes for this case are mixed structured/unstructured (hybrid) topologies. The mesh in the immediate vicinity of the wing uses a quad-dominant approach to produce mostly hexahedra in the wing boundary layer. This has most of the advantages of an unstructured triangular mesh in terms of ease of meshing and face isotropy in the interior, but has fewer elements for a comparable node count. A potential disadvantage is that there is no way to produce a mixed hex/tet mesh without the introduction of pyramid elements, which can cause convergence and accuracy problems. There is also a refined region around the tip inside the wing box to ensure resolution of the formation of the wing tip vortex.

Further downstream, there is a fully structured hex mesh, expanding slightly and covering the path of the tip vortex downstream as measured in the experiment.

The balance of the test section mesh is unstructured tets (except as noted below), while another structured block is used upstream of the test section.

The meshes first produced used the Discontinuous Galerkin (DG) non-conformal interface between the hexahedral tip vortex mesh and the fully unstructured test section mesh. Due to the relative novelty of the DG approach and our lack of familiarity with its performance in Nalu-Wind, it was decided that a more conservative traditional, conformal

interface between the blocks was preferable. Therefore, the tetrahedral test section block interfaces to the hexahedral tip vortex block and the upstream block using node-matched pyramid elements.

Notes on Geometry

- The trailing edge geometry of the NACA0018 airfoil isn't given in either the McAlister report nor in the original NACA publications describing it. Therefore, for ease of meshing, a rounded trailing edge was used.
- In order to capture at least the gross blockage effects, the model support structure in the wind tunnel is modeled, and the tunnel walls are at the correct locations. However, in an effort to keep the mesh size low, the tunnel walls and the support are modeled as slip walls and not as viscous.
- Most of the McAlister cases of interest were performed using a square wingtip. The initial mesh, however, uses the rounded tip described in McAlister. We will eventually produce a square tip mesh as this is both more interesting and has more-complete results.

Surface Mesh




Fig. 1.1: The surface mesh near the tip, as viewed from above.

Statistics of Current Mesh (`grid07_conformal10.exo`)

Node count: 58M

Element Count: 192M total, consisting of:

- 158M tets




Fig. 1.2: A close-up view of the tip and trailing edge, showing rounded tip and trailing edge and quad-dominant surface mesh.

- 2.5M pyramids
- 1.1M wedges
- 30M hexes

Max. Centroid Skew: 0.866; $52 > 0.8$

Max. Included Angle: 177 degrees; $7 > 170$ degrees

Max. Volume Ratio: 22; $12 > 20$

Max. Aspect Ratio: 346

Wall Spacing on Wing: 8.8×10^{-5} m

T-Rex Growth Rate: 1.2

Full/Max Layers in Tip Block: 19/19 (limited to preserve quality)

Full/Max Layers in Wing Block: 19/33

Lessons Learned

- We need *a lot* of resolution to resolve and advect the tip vortex: on the order of 2–3mm edge length.
- Due to the mathematics of physical space, small changes in the maximum edge length in a block lead to large changes in the final mesh size. For example, changing the maximum edge length from 0.0025m to 0.003m produces nearly a factor of 2 difference in the element count in the isotropic portion of the mesh: $(0.003/0.0025)^3 \approx 1.73$.

- Heuristically, volume ratios should ideally be < 20 . Slightly larger volume ratios are acceptable as long as there are no steep gradients passing through these elements.
- Aspect ratios should be $< 1000:1$
- Centroid skewness is a better measure than the other skewness metrics as it is more even across element topologies
 - equiangle skewness is also OK, but is stricter and can give misleadingly high readings for some tets
 - equivolume skewness is useless for tets
- Centroid skewness should be < 0.8 ; however, skewness as high as the low 0.9s (usually associated with topology transitions) is acceptable as long as:
 - the skewed cells are far away from large gradients; and
 - there are no more than a handful.

General Pointwise Tips

- Maximum aspect ratio for quads in domains should be ≤ 4 for good quality extrusions.
- Maximum included angle should be ≤ 170 degrees. The usual exceptions for regions with small gradients *should* apply here, but there may be additional restrictions due to the elliptic nature of the incompressible flow equations.
- It can be beneficial to push poor quality cells out of the boundary layer by increasing the minimum number of T-Rex layers.
- One can set the maximum number of layers to prevent different numbers of layers in a block and its adjacent domains. This can eliminate some poor-quality tetrahedra.

References

1.3 Nalu-Wind - Examples

The example cases presented are meant to get users started using nalu-wind.

1.3.1 Introduction

A collection of examples for running large-eddy simulations of wind plant aerodynamics is available under `nalu-wind/examples/`. A `setup.yaml` file is included with all cases. This file has the most relevant parameters that can be modified for each case of interest.

A set of Python utilities is included with the examples. These utilities are meant to simplify the process of generating input files for running Nalu and plot results from the simulations.

Each case has a default collection of template input files in `./template_input_files`. The executable **nalu_input_fileX** will take the input files and modify them according to the inputs set in `setup.yaml` and generate new and ready to use input files.

Instructions to compile **Nalu** are provided in [Building Nalu-Wind](#). The `wind-utils` repository needs to be compiled to have access to all the pre-processing utilities used in the example cases. The **wind-utils** repository can be downloaded as a submodule by running this command inside the `nalu-wind/` repository:

```
git submodule init && git submodule update
```

Now, **wind-utils** can be compiled with **Nalu** by enabling the compilation flag:

```
-DENABLE_WIND_UTILS=ON
```

during CMake configure phase. Subsequent make install will install all **wind-utils** executables along side **nalux** under the same installation prefix.

The general instructions to run each case

1. Modify the simulation parameters in the `setup.yaml` file.
2. Execute the **nalu_input_fileX** script with the `setup.yaml` file as an input.
3. Generate the mesh using **abl_mesh** from nalu wind utils.
4. Generate the initial condition using **nalu_preprocess**.
5. Run the simulation using **nalux**.

Setting up the environment

In order to use the Python utilities to create the input files and post-process some of the data, a proper environment needs to be set. The user can add these libraries to their Python environment, or use conda to create the environment needed. Instruction to install conda can be found [here](#).

The new environment can be created through conda using:

```
conda create -n nalu_python -c conda-forge python=3.6 numpy ruamel.yaml_
↪netCDF4 matplotlib scipy pandas
```

This new environment will allow the execution of **nalu_input_fileX**. The environment is saved in THE USER system, so it needs to be created only once. After that, it just needs to be activated.

Now, to use the environment run:

```
source activate nalu_python
```

The **nalu_input_fileX** script

This code is an executable which takes as an input a set-up file. The executable will read in the set-up file, and create a new nalu input file based on the parameters specified. Executing the code with the `-h` flag will provide the necessary information:

```
./nalu_input_fileX -h
```

The **setup.yaml** file

This file includes the inputs to be modified for a case. This example is for a Neutral Atmospheric Boundary Layer simulation.

```
#####
↪###
#   This is the input file for an ABL Simulation.
#   All the parameters in this file will be used to create Nalu inputs.
#####
↪###

#####
# These are the example input files and the files which will be generated.
#####
# The old input files which will be modified
template_input: template_input_files/ablNeutralEdge.yaml
template_preprocess: template_input_files/nalu_preprocess.yaml

# The name of the new input files
new_input: abl_simulation.yaml
new_preprocess: abl_preprocess.yaml

# Establish if the simulation is a restart or not (yes/no)
restart: no

#####
# Wind speed and temperature profiles
#####
# Wind speed at specified height [m/s].
U0Mag: 8.0
# Wind direction [deg]. A direction of 270 deg means the wind is coming from_
↪the
#   west, which is from left to right.
#
#           N 0deg
#           |
#           |
#           |
#   W 270deg --- E 90deg
#           |
#           |
#           |
#           S 180deg
#
wind_dir: 270.0
# Height at which to drive mean wind to U0Mag/dir [m].
wind_height: 90.0
# Temperature values [K] at each height [m] for initial condition
temperature_heights: [ 0, 650.0, 750.0, 1000.0 ]
temperature_values: [300.0, 300.0, 308.0, 308.75]

#####
# Material properties
#####
# Density [kg/m^3]
density: 1.
# Kinematic viscosity [m^2/s].
nu: 1.0E-5
# Reference potential temperature [K].
TRef: 300.0
# Latitude on the Earth of the site [deg].
```

(continues on next page)

(continued from previous page)

```

latitude:                41.3

#####
# Bottom wall
#####
# Wall-normal component of temperature flux at wall.
# A negative value is flux into domain [K-m/s].
qwall:                    0.
# Surface roughness (m).
z0:                        0.15

#####
# Time controls
#####
# Time-step [s]
time_step: 0.5
# Total number of time to run [s]
total_run_time: 20000
# Check for CFL condition

#####
# Mesh
#####
# Here are the mesh properties
mesh:
# Generate the mesh or read from an input file (yes/no)
generate: yes
# The name of the mesh file
# If the mesh is generated the preprocessor will generate a mesh
# If not the code will read in this mesh
mesh_file: mesh_abl.exo
# The domain bounds [m]
domain_bounds_x: [0.0, 5000.0]
domain_bounds_y: [0.0, 5000.0]
domain_bounds_z: [0.0, 1000.0]

# The number of grid points in the x, y, z coordinates
# Change this variable
number_of_cells: [500, 500, 100]

#####
# Output
#####
# How often to write output [s]
# This is time-steps in nalu input
output_frequency: 10000
# Output file name. It will create the directory
# Change .e
output_data_base_name: output/abl_5km_5km_1km_neutral.exo

# These are the plane averaged statistics
boundary_layer_statistics:
# The file to write output (netcdf)
stats_output_file: abl_statistics.nc
# How often sample the statistics
time_hist_output_frequency: 1

```

(continues on next page)

(continued from previous page)

```
#####
# Add restart capability
# Default to write last time-step
#####

#####
# Upper wall boundary condition
# This is an optional flag
# If not specified, the gradient will be computed
#   based on the initial condition
#   at the top of the domain
#####
# Potential temperature gradient above the strong inversion (K/m).
# Compute this from the last points
# TGradUpper:          0.003
```

1.3.2 Peregrine

Here are the instructions to use Nalu and these examples on NREL's HPC system Peregrine.

Initial Setup

These steps need to be completed only once to setup the appropriate Nalu environment on the Peregrine system.

1. Create the conda environment:

```
module load conda
conda create -n nalu_python -c conda-forge python=3.6 numpy ruamel.yaml netCDF4_
↪matplotlib scipy pandas
```

2. Set the Nalu environment. This is the same environment used to compile Nalu. If this has not been set, then it can be done by adding the following function to `${HOME}/.bash_profile`:

```
function nalu_env {
  module purge
  # Load the python environment
  module load conda
  source activate nalu_python

  local mod_dir=/nopt/nrel/ecm/ecp/base/modules/
  module use ${mod_dir}/gcc-6.2.0
  module load gcc/6.2.0

  compiler=${1:-gcc}

  case ${compiler} in
    gcc)
      module unuse ${mod_dir}/intel-18.1.163
      module load binutils openmpi/3.1.1 netlib-lapack cmake
      ;;
    intel)
      module load intel-parallel-studio/cluster.2018.1
      module use ${mod_dir}/intel-18.1.163
      module load binutils intel-mpi intel-mkl cmake
```

(continues on next page)

(continued from previous page)

```
        ;;  
    esac  
}
```

Source the new `${HOME}/.bash_profile`:

```
source ${HOME}/.bash_profile
```

3. Clone the Nalu repository by:

```
git clone https://github.com/Exawind/nalu-wind.git
```

Now the environment is ready to run the examples. To get started, go to: [Atmospheric Boundary Layer - Neutral](#).

Running Every Case

Every time the user logs into Peregrine and wants to run a case, these steps need to be completed:

1. Load the nalu environment:

```
nalu_env
```

2. Copy the executables from the public location to the directory where the case is:

```
cp /projects/windsim/nalu-wind-executables/* .
```

The system is now ready to compile and use Nalu.

1.3.3 Atmospheric Boundary Layer - Neutral

This case is a large-eddy simulation of a neutral atmospheric boundary layer. The case uses periodic boundary conditions on the sides (east, west, north, south), a wall-model in the bottom wall and a stress free boundary condition at the top of the domain. A proportional controller is used to drive the velocity at a given height. The controller adjusts the forcing at each time-step to match a given planar average velocity at a given height. More information about the controller can be found in [ABL Forcing Source Terms](#). It takes about 10,000 [s] for the turbulence to develop. The example runs for 20,000 [s].

Step by step instructions to run the case

1. Load the appropriate Nalu environment. This requires loading the libraries and Python environment as described in [Setting up the environment](#). For users on Peregrine the function defined in [Initial Setup](#) should suffice:

```
nalu_env
```

2. Go to the directory where the case is:

```
cd nalu-wind/examples/abl_neutral/
```

3. Modify the `setup.yaml` file to include all the necessary simulation parameters.
4. Run the executable and provide the `setup.yaml` file as input:


```
../nalu_input_fileX -s setup.yaml
```

For users on Peregrine, now copy the executables to the case directory:

```
cp /projects/windsim/nalu-wind-executables/* .
```

5. Generate the mesh:

```
./abl_mesh -i abl_preprocess.yaml
```

6. Generate the initial condition:

```
./nalu_preprocess -i abl_preprocess.yaml
```

7. Run the nalu executable:

```
mpirun -np 600 naluX -i abl_simulation.yaml
```

In this example 600 processors are used, but any number of processors could be used. Target 50K elements per core for choosing number of MPI cores.

Post-processing

Nalu's output is generated at runtime. The `abl_plots.py` Python script is used to plot simulation results. The script will load the plane-averaged statistics and plot them as function of time and height. To run the script, load the Python environment if needed, and run the Python script:

```
python abl_plots.py
```

1.3.4 Actuator Line Model in Uniform Inflow

This case is a large-eddy simulation of 2 aligned wind turbines under uniform inflow. The turbines are represented using an actuator line model. The turbine model is the NREL 5MW Reference. The case has uniform inflow boundary condition with an outflow boundary condition on the other end. The sides are specified as zero stress boundary conditions. The wind turbine aerodynamic forces are computed using OpenFAST.

Step by step instructions to run the case

1. Load the appropriate Nalu environment. This requires loading the libraries and Python environment as described in *Setting up the environment*. For users on Peregrine the function defined in *Initial Setup* should suffice:

```
nalu_env
```

2. Go to the directory where the case is:

```
cd nalu-wind/examples/turbine_uniform_inflow/
```

3. Modify the `setup.yaml` file to include all the necessary simulation parameters.
4. Run the executable and provide the `setup.yaml` file as input:

```
../nalu_input_fileX -s setup.yaml
```

For users on Peregrine, now copy the executables to the case directory:

```
cp /projects/windsim/nalu-wind-executables/* .
```

5. Generate the mesh:

```
./abl_mesh -i alm_preprocess.yaml
```

6. Run the nalu executable:

```
mpirun -np 24 naluX -i alm_simulation.yaml
```

Post-processing

The turbine output is generated at runtime. The `plot_alm.py` Python script is used to plot turbine output. The script will load the OpenFAST data and plot it as a function of time. To run the script, load the Python environment if needed, and run the Python script:

```
python plot_alm.py
```

1.3.5 Wind Farm

This case is a large-eddy simulation of a wind farm under neutral stability conditions. The simulation requires 3 steps: 1) running a precursor atmospheric boundary layer simulation, 2) sampling the data from the boundaries, and 3) running the wind farm simulation with inflow/outflow boundary conditions.

This examples guides the user through all these steps.

Case 1: Precursor simulation

This example follows the same procedure as *Atmospheric Boundary Layer - Neutral*. The only difference is that now, the boundary data is sampled.

1. Load the appropriate Nalu environment. This requires loading the libraries and Python environment as described in *Setting up the environment*. For users on Peregrine the function defined in *Initial Setup* should suffice:

```
nalu_env
```

2. Go to the directory where the case is:

```
cd nalu-wind/examples/wind_farm/case1/
```

3. Modify the `setup.yaml` file to include all the necessary simulation parameters.
4. Run the executable and provide the `setup.yaml` file as input:

```
../../nalu_input_fileX -s setup.yaml
```

5. Generate the mesh:

```
./abl_mesh -i case_1_precursor_preprocess.yaml
```

6. Generate the initial condition:

```
../nalu_preprocess -i case_1_precursor_preprocess.yaml
```

7. Run the nalu executable:

```
mpirun -np 8 ../nalux -i case_1_precursor_simulation.yaml
```

In this example 8 processors are used, but any number of processors could be used. Target 50K elements per core for choosing number of MPI cores.

Case 2: Sampling boundary data

Now, a simulation for sampling the boundary data is performed. This is the same as the previous simulation, but with 2 changes: 1) Data sampling from the boundary conditions. 2) Different time interval

1. Load the appropriate Nalu environment. This requires loading the libraries and Python environment as described in *Setting up the environment*. For users on Peregrine the function defined in *Initial Setup* should suffice:

```
nalu_env
```

2. Go to the directory where the case is:

```
cd nalu-wind/examples/wind_farm/case2
```

3. Modify the `setup.yaml` file to include all the necessary simulation parameters.
4. Run the executable and provide the `case_2_setup_abl_precursor_boundary_data.yaml` file as input:

```
../..//nalu_input_fileX -s setup.yaml
```

5. Generate the boundary data to sample:

```
../nalu_preprocess -i case_2_boundary_data_preprocess.yaml
```

6. Run the nalu executable:

```
mpirun -np 8 ../nalux -i case_2_abl_precursor_boundary_data.yaml
```

In this example 8 processors are used, but any number of processors could be used. Target 50K elements per core for choosing number of MPI cores.

Case 3: Wind farm simulation

The last part is the wind farm simulation. In this example, the simulation is run using the same resolution as the precursor. However, it is possible to perform local refinement near the turbines. Full instructions are provided in *wind-utils*.

1. Load the appropriate Nalu environment. This requires loading the libraries and Python environment as described in *Setting up the environment*. For users on Peregrine the function defined in *Initial Setup* should suffice:

```
nalu_env
```

2. Go to the directory where the case is:

```
cd nalu-wind/examples/wind_farm/case3
```

3. Modify the `setup.yaml` file to include all the necessary simulation parameters.
4. Run the executable and provide the `setup.yaml` file as input:

```
../../nalu_input_fileX -s setup.yaml
```

5. Run the nalu executable:

```
mpirun -np 8 ../naluX -i case_3_wind_farm.yaml
```

In this example 8 processors are used, but any number of processors could be used. Target 50K elements per core for choosing number of MPI cores.

Post-processing

The data generated by this example can be post-processed using the same scripts provided in *Atmospheric Boundary Layer - Neutral* and *Actuator Line Model in Uniform Inflow*.

DEVELOPER MANUAL

2.1 Testing Nalu-Wind

Nalu-Wind's regression tests and unit tests are run nightly using the GCC and Intel compilers against the Trilinos master and development branches on a machine at NREL. The results can be seen at the [CDash Nalu-Wind website](#).

2.1.1 Running Tests Locally

The nightly tests are implemented using CTest and these same tests are available to developers to run locally as well. Due to the nature of error propagation of calculations in computers, results of regression testing can be difficult to keep consistent. Output from Nalu-Wind can vary from established reference data for the regression tests based on the compiler you are using, the types of optimizations set, and the versions of the third-party libraries Nalu-Wind utilizes, along with the blas/lapack implementation in use. Therefore it may make sense when you checkout Nalu-Wind to create your own reference data for the tests for the machine and configuration you are using, which is described later in this document. Or you can use a lower tolerance when running the tests. At the moment, a single tolerance is chosen in which to use for all the tests. The following instructions will describe how to run Nalu-Wind's tests.

Since Nalu-Wind's tests require a large amount of data (meshes), this data is hosted in a separate repository from Nalu-Wind. This mesh repo is set as a submodule in the `reg_tests/mesh` directory in the main Nalu-Wind repository. Submodule repos are not checked out by default, so either use `git submodule init` and then `git submodule update` to clone it in your checkout of Nalu-Wind, or when you first clone Nalu-Wind you can also use `git clone --recursive <repo_url>` to checkout all submodules as well.

Once this submodule is initialized and cloned, you will need to configure Nalu-Wind with testing on. To configure Nalu-Wind with testing enabled, in Nalu-Wind's existing `build` directory, we will run this command:

```
cmake -DTrilinos_DIR:PATH=`spack location -i nalu-trilinos` \  
      -DYAML_DIR:PATH=`spack location -i yaml-cpp` \  
      -DENABLE_TESTS:BOOL=ON \  
      ..
```

Note we have chosen to originally build Nalu-Wind with Spack in this case, hence the use of `spack location -i <package>` to locate our YamI and Trilinos installations. Then we use `-DENABLE_TESTS:BOOL=ON` to enable CTest. Once Nalu-Wind is configured, you should be able to run the tests by building Nalu-Wind in the `build` directory, and running `make test` or `ctest`. Looking at `ctest -h` will show you many ways you can run tests and choose which tests to run.

There are advantages to using CTest, such as being able to run subsets of the tests, or tests matching a particular regular expression for example. To do so, in the `build` directory, you can run `ctest -R femHC` to run the test matching the `femHC` regular expression. Other useful capabilities are `ctest --output-on-failure` to see test outputs when they fail, `ctest --rerun-failed` to only run the tests that previously failed, `ctest --print-labels` to see the test labels, and `ctest -L unit` to run the tests with label 'unit' for example. All testing related log files

and output can be seen in `Nalu-Wind/build/Testing/Temporary` and `Nalu-Wind/build/reg_tests` after the test have been run.

To define your own tolerance for tests, at configure time, add `-DTEST_TOLERANCE=0.0001` for example to the Nalu-Wind CMake configure line.

Updating Reference Data for Your Machine

When running the tests, the norms for each test are the output and they are ‘diffed’ against the ‘gold’ norms that we have established for each test. To dictate whether or not a test passes, we use a chosen tolerance in which we allow the results to deviate from the ‘gold’ norm. As stated earlier, these ‘gold’ norms are not able to reflect every configuration of Nalu-Wind, per compiler, optimization, TPL versions, blas/lapack version, etc. This tolerance is currently defined in the `CMakeLists.txt` in Nalu-Wind’s `reg_tests` directory. This tolerance can also be passed into Nalu-Wind at configure time using `-DTEST_TOLERANCE=0.0000001` for example. To update the ‘gold’ norms locally to your configuration, merely run the tests once, and copy the `*.norm` files in the `build/reg_tests/test_files` directory to the corresponding test location in `reg_tests/test_files` while overwriting the current ‘gold’ norms.

In regards to ‘official’ gold norms, Linux with GCC 6.4.0, netlib-blas/lapack 3.8.0, and the following TPL versions are officially tested:

```
openmpi@1.10.4
boost@1.66.0
cmake@3.9.4
parallel-netcdf@1.8.0
yaml-cpp@develop
hdf5@1.10.1
netcdf@4.4.1.1
zlib@1.2.11
superlu@4.3
```

2.1.2 Adding Tests to Nalu-Wind

The testing infrastructure is almost completely confined to the `reg_tests` directory. To add a test to Nalu-Wind, we need to add the test name, and create a test directory to place the input files and gold norms for the test. First, the test itself can be added to the list of CTest tests by adding a line to the `CTestList.cmake` file. For a single regression test, provided it is similar to the categories shown at the top of the `CTestList.cmake` file, it can likely be added with a single line using the test name and amount of processes you would like to run the test with and choosing the correct function to use. For example:

```
add_test_r(mytest 6)
```

After this has been done, in the `reg_tests/test_files` directory, you should add a directory corresponding to your test name and include the input file, `mytest.i`, and reference output file `mytest.norm.gold`. If you are using an xml file that doesn’t exist in the `xml` directory, you will need to commit that as well.

To see commands used when running the tests, see the functions at the top of the `CTestList.cmake` file. These functions ultimately create `CTestTestFile.cmake` files in the CMake build directory at configure time. You can see the exact commands used for each test after configure in the `build/reg_tests/CTestTestFile.cmake` file.

Note if your test doesn’t conform to an existing archetype, a new function in `CTestList.cmake` may need to be created. Also, if you are using a mesh file that doesn’t exist in the mesh repo, you will need to add it, and update the submodule in the Nalu-Wind main repo to use the latest commit of the mesh submodule repo.

2.1.3 Adding Testing Machines to CDash

To add a testing machine that will post results to CDash first means that you should have all software dependencies satisfied for Nalu-Wind. Next the script located at [CTestNightlyScript.cmake](#) can be run for example as:

```
ctest \
  -DNIGHTLY_DIR=${NALU_WIND_TESTING_DIR} \
  -DYAML_DIR=${YAML_INSTALL_DIR} \
  -DTRILINOS_DIR=${TRILINOS_INSTALL_DIR} \
  -DHOST_NAME=machine.domain.com \
  -DEXTRA_BUILD_NAME=Linux-gcc-whatever \
  -VV -S ${NALU_WIND_DIR}/reg_tests/CTestNightlyScript.cmake
```

In this case `${NALU_WIND_TESTING_DIR}` is one directory above where the Nalu-Wind repo has been checked out. This runs CTest in scripting mode with verbosity on and it will update the Nalu-Wind repo with the latest revisions, configure, build, test, and finally submit results to the CDash site. Since CTest does the building, it needs to know the locations of Yaml and Trilinos. For examples of nightly testing, refer to the testing scripts currently being run [here](#).

2.2 Source Code Documentation

The source documentation is extracted from the C++ files using Doxygen.

2.2.1 Simulation – Nalu Top-level Interface

class Simulation

Realms

Realm is a Nalu abstraction of a set of equations that are solved on a computational domain, represented by an Exodus-II mesh. A simulation can contain multiple Realms and that can interact via `sierra::nalu::Transfer` instance. `InputOutputRealm` is a special type of Realm that exists solely to provide data (input) or extract a subset of data from another `Realm`.

class Realm

Representation of a computational domain and physics equations solved on this domain.

Subclassed by `sierra::nalu::InputOutputRealm`

Public Functions

void **set_hypre_global_id**()
Initialize the HYPRE global row IDs.

See `Realm::hypreGlobalId_`

void **check_job**(bool *get_node_count*)
check job for fitting in memory

Public Members

stk::mesh::PartVector **bcPartVec_**

Vector holding side sets that have been registered with the boundary conditions in the input file.

The member is intended to for use in `Realm::enforce_bc_on_exposed_faces` to check for “exposed surfaces” that might have not been assigned BCs in the input file.

stk::mesh::EntityId **hypreILower_**

The starting index (global) of the HYPRE linear system in this MPI rank.

Note that this is actually the offset into the linear system. This index must be adjusted accordingly to account for multiple degrees of freedom on a particular node. This is performed in *sierra::nalu::HypreLinearSystem*.

stk::mesh::EntityId **hypreIUpper_**

The ending index (global) of the HYPRE linear system in this MPI rank.

Note that this is actually the offset into the linear system. This index must be adjusted accordingly to account for multiple degrees of freedom on a particular node. This is performed in *sierra::nalu::HypreLinearSystem*.

stk::mesh::EntityId **hypreNumNodes_**

The total number of HYPRE nodes in the linear system.

Note that this is not an MPI rank local quantity

HypreIDFieldType ***hypreGlobalId_** = {nullptr}

Global Row IDs for the HYPRE linear system.

The HYPRE IDs are different from STK IDs and `Realm::naluGlobalId_` because HYPRE expects contiguous IDs for matrix rows and further requires that the IDs be ordered across MPI ranks; i.e., $\text{startIdx}(\text{MPI_rank} + 1) = \text{endIdx}(\text{MPI_rank}) + 1$.

bool **hypreIsActive_** = {false}

Flag indicating whether Hypre solver is being used for any of the equation systems.

class InputOutputRealm : public sierra::nalu::Realm

class Realms

Time Integration

class TimeIntegrator

Linear Solver Interface

class LinearSystem

Subclassed by *sierra::nalu::HypreLinearSystem*, *sierra::nalu::TpetraLinearSystem*,
sierra::nalu::TpetraSegregatedLinearSystem

Public Functions

virtual void **buildDirichletNodeGraph** (const stk::mesh::PartVector&)

Process nodes that belong to Dirichlet-type BC.

virtual void buildDirichletNodeGraph (**const** std::vector<stk::mesh::Entity>&)

Process nodes as belonging to a Dirichlet-type row.

See the documentation/implementation of `sierra::nalu::FixPressureAtNodeAlgorithm` for an example of this use case.

virtual void resetRows (**const** std::vector<stk::mesh::Entity> &*nodeList*, **const** unsigned *beginPos*, **const** unsigned *endPos*, **const** double *diag_value* = 0.0, **const** double *rhs_residual* = 0.0) = 0

Reset LHS and RHS for the given set of nodes to 0.

Parameters

- *nodeList*: A list of STK node entities whose rows are zeroed out
- *beginPos*: Starting index (usually 0)
- *endPos*: Terminating index (1 for scalar quantities; nDim for vectors)

class LinearSolver

An abstract representation of a linear solver in Nalu.

Defines the basic API supported by the linear solvers for use within Nalu. See concrete implementations such as `sierra::nalu::TpetraLinearSolver` for more details.

Subclassed by `sierra::nalu::HyprDirectSolver`, `sierra::nalu::TpetraLinearSolver`

Public Functions

virtual PetraType getType () = 0

Type of solver instance as defined in `sierra::nalu::PetraType`.

virtual void destroyLinearSolver () = 0

Utility method to cleanup solvers during simulation.

bool &recomputePreconditioner ()

Flag indicating whether the preconditioner is recomputed on each invocation.

bool &reusePreconditioner ()

Flag indicating whether the preconditioner is reused on each invocation.

void zero_timer_precond ()

Reset the preconditioner timer to 0.0 for future accumulation.

double get_timer_precond ()

Get the preconditioner timer for the last invocation.

bool &activeMueLu ()

Flag indicating whether the user has activated MueLU.

LinearSolverConfig ***getConfig** ()

Get the solver configuration specified in the input file.

Public Members

std::string **name_**

User-friendly identifier for this particular solver instance.

class TpetraLinearSystem : **public** `sierra::nalu::LinearSystem`

Public Functions

```
virtual void resetRows (const std::vector<stk::mesh::Entity> &nodeList, const unsigned begin-  
Pos, const unsigned endPos, const double diag_value = 0.0, const  
double rhs_residual = 0.0)  
Reset LHS and RHS for the given set of nodes to 0.
```

Parameters

- nodeList: A list of STK node entities whose rows are zeroed out
- beginPos: Starting index (usually 0)
- endPos: Terminating index (1 for scalar quantities; nDim for vectors)

Transfers

```
class Transfer
```

```
class Transfers
```

2.2.2 Equation Systems

```
class EquationSystem
```

Base class representation of a PDE.

EquationSystem defines the API supported by all concrete implementations of PDEs for performing the following actions:

- Register computational fields
- Register computational algorithms for interior domain and boundary conditions
- Manage solve and update of the PDE for a given timestep

Subclassed by *sierra::nalu::ContinuityEquationSystem*, *sierra::nalu::EnthalpyEquationSystem*,
sierra::nalu::HeatCondEquationSystem, *sierra::nalu::LowMachEquationSystem*,
sierra::nalu::MassFractionEquationSystem, *sierra::nalu::MatrixFreeHeatCondEquationSystem*,
sierra::nalu::MatrixFreeLowMachEquationSystem, *sierra::nalu::MeshDisplacementEquationSystem*,
sierra::nalu::MixtureFractionEquationSystem, *sierra::nalu::MomentumEquationSystem*,
sierra::nalu::ProjectedNodalGradientEquationSystem, *sierra::nalu::ShearStressTransportEquationSystem*,
sierra::nalu::SpecificDissipationRateEquationSystem, *sierra::nalu::TurbKineticEnergyEquationSystem*,
sierra::nalu::WallDistEquationSystem

Public Functions

```
virtual void solve_and_update ()
```

Assemble the LHS and RHS and perform linear solve for prescribed number of iterations.

This method is invoked in *EquationSystems::solve_and_update* method as shown below

```
pre_iter_work();  
// Iterate over all equation systems  
for (auto eqsys: equationSystems_) {  
    eqsys->pre_iter_work();  
}
```

(continues on next page)

(continued from previous page)

```

eqsys->solve_and_update();           //<<<< Assemble and solve system
eqsys->post_iter_work();
}
post_iter_work();

```

See *EquationSystems::solve_and_update*

virtual void pre_iter_work()

Perform setup tasks before entering the solve and update step.

This method is invoked in *EquationSystems::solve_and_update* method as shown below

```

pre_iter_work();
// Iterate over all equation systems
for (auto eqsys: equationSystems_) {
    eqsys->pre_iter_work();           //<<<< Pre-iteration setup
    eqsys->solve_and_update();
    eqsys->post_iter_work();
}
post_iter_work();

```

See *EquationSystems::solve_and_update*

virtual void post_iter_work()

Perform setup tasks after the solve and update step.

This method is invoked in *EquationSystems::solve_and_update* method as shown below

```

pre_iter_work();
// Iterate over all equation systems
for (auto eqsys: equationSystems_) {
    eqsys->pre_iter_work();
    eqsys->solve_and_update();
    eqsys->post_iter_work();         //<<<< Post-iteration actions
}
post_iter_work();

```

See *EquationSystems::solve_and_update*

virtual void post_iter_work_dep()

Deprecated post iteration work logic.

template<typename **T** = DoubleType>

PecletFunction<**T**> ***ngp_create_peclet_function**(const std::string &dofName)

Create and return an instance of PecletFunction on device for use with Kernel.

virtual void solution_update(const double *delta_frac*, const stk::mesh::FieldBase &*delta*,
const double *field_frac*, stk::mesh::FieldBase &*field*, const
unsigned *numComponents* = 1, const stk::topology::rank_t *rank*
= stk::topology::NODE_RANK)

Update field with delta solution of linear solve.

Public Members

`std::vector<PecletFunctionBase*> ngpPecletFunctions_`
Track NGP instances of PecletFunction.

`std::vector<AlgorithmDriver*> preIterAlgDriver_`
List of tasks to be performed before each *EquationSystem::solve_and_update*.

`std::vector<AlgorithmDriver*> postIterAlgDriver_`
List of tasks to be performed after each *EquationSystem::solve_and_update*.

`size_t linsysWriteCounter_ = {0}`
Counter to track the number of linear system outputs.

Move this to *EquationSystem* instead of *LinearSystem* so that we don't reset the counter when performing matrix reinitializations.

class LowMachEquationSystem: public *sierra::nalu::EquationSystem*
Low-Mach formulation of the Navier-Stokes Equations.

This class is a thin-wrapper around *sierra::nalu::ContinuityEquationSystem* and *sierra::nalu::MomentumEquationSystem* that orchestrates the interactions between the velocity and the pressure Poisson solves in the *LowMachEquationSystem::solve_and_update* method.

Public Functions

virtual void pre_iter_work()
Perform setup tasks before entering the solve and update step.

This method is invoked in *EquationSystems::solve_and_update* method as shown below

```
pre_iter_work();  
// Iterate over all equation systems  
for (auto eqsys: equationSystems_) {  
    eqsys->pre_iter_work();           //<<<< Pre-iteration setup  
    eqsys->solve_and_update();  
    eqsys->post_iter_work();  
}  
post_iter_work();
```

See *EquationSystems::solve_and_update*

virtual void solve_and_update()
Assemble the LHS and RHS and perform linear solve for prescribed number of iterations.

This method is invoked in *EquationSystems::solve_and_update* method as shown below

```
pre_iter_work();  
// Iterate over all equation systems  
for (auto eqsys: equationSystems_) {  
    eqsys->pre_iter_work();  
    eqsys->solve_and_update();           //<<<< Assemble and solve system  
    eqsys->post_iter_work();  
}  
post_iter_work();
```

See *EquationSystems::solve_and_update*

```
class EnthalpyEquationSystem : public sierra::nalu::EquationSystem
```

Public Functions

```
void solve_and_update()
```

Assemble the LHS and RHS and perform linear solve for prescribed number of iterations.

This method is invoked in [EquationSystems::solve_and_update](#) method as shown below

```
pre_iter_work();
// Iterate over all equation systems
for (auto eqsys: equationSystems_) {
    eqsys->pre_iter_work();
    eqsys->solve_and_update();           //<<<< Assemble and solve system
    eqsys->post_iter_work();
}
post_iter_work();
```

See [EquationSystems::solve_and_update](#)

```
void post_iter_work_dep()
```

Deprecated post iteration work logic.

```
class TurbKineticEnergyEquationSystem : public sierra::nalu::EquationSystem
```

Public Functions

```
void solve_and_update()
```

Assemble the LHS and RHS and perform linear solve for prescribed number of iterations.

This method is invoked in [EquationSystems::solve_and_update](#) method as shown below

```
pre_iter_work();
// Iterate over all equation systems
for (auto eqsys: equationSystems_) {
    eqsys->pre_iter_work();
    eqsys->solve_and_update();           //<<<< Assemble and solve system
    eqsys->post_iter_work();
}
post_iter_work();
```

See [EquationSystems::solve_and_update](#)

```
class ShearStressTransportEquationSystem : public sierra::nalu::EquationSystem
```

Public Functions

```
virtual void solve_and_update()
```

Assemble the LHS and RHS and perform linear solve for prescribed number of iterations.

This method is invoked in [EquationSystems::solve_and_update](#) method as shown below

```
pre_iter_work();
// Iterate over all equation systems
for (auto eqsys: equationSystems_) {
    eqsys->pre_iter_work();
    eqsys->solve_and_update();           //<<<< Assemble and solve system
    eqsys->post_iter_work();
}
post_iter_work();
```

See *EquationSystems::solve_and_update*

virtual void post_iter_work()

Perform setup tasks after the solve and update step.

This method is invoked in *EquationSystems::solve_and_update* method as shown below

```
pre_iter_work();
// Iterate over all equation systems
for (auto eqsys: equationSystems_) {
    eqsys->pre_iter_work();
    eqsys->solve_and_update();
    eqsys->post_iter_work();           //<<<< Post-iteration actions
}
post_iter_work();
```

See *EquationSystems::solve_and_update*

class HeatCondEquationSystem: public sierra::nalu::EquationSystem

Public Functions

void solve_and_update()

Assemble the LHS and RHS and perform linear solve for prescribed number of iterations.

This method is invoked in *EquationSystems::solve_and_update* method as shown below

```
pre_iter_work();
// Iterate over all equation systems
for (auto eqsys: equationSystems_) {
    eqsys->pre_iter_work();
    eqsys->solve_and_update();           //<<<< Assemble and solve system
    eqsys->post_iter_work();
}
post_iter_work();
```

See *EquationSystems::solve_and_update*

class MassFractionEquationSystem: public sierra::nalu::EquationSystem

Public Functions

void solve_and_update()

Assemble the LHS and RHS and perform linear solve for prescribed number of iterations.

This method is invoked in *EquationSystems::solve_and_update* method as shown below

```
pre_iter_work();
// Iterate over all equation systems
for (auto eqsys: equationSystems_) {
    eqsys->pre_iter_work();
    eqsys->solve_and_update();           //<<<< Assemble and solve system
    eqsys->post_iter_work();
}
post_iter_work();
```

See *EquationSystems::solve_and_update*

```
class MixtureFractionEquationSystem: public sierra::nalu::EquationSystem
```

Public Functions

```
void solve_and_update ()
```

Assemble the LHS and RHS and perform linear solve for prescribed number of iterations.

This method is invoked in *EquationSystems::solve_and_update* method as shown below

```
pre_iter_work();
// Iterate over all equation systems
for (auto eqsys: equationSystems_) {
    eqsys->pre_iter_work();
    eqsys->solve_and_update();           //<<<< Assemble and solve system
    eqsys->post_iter_work();
}
post_iter_work();
```

See *EquationSystems::solve_and_update*

```
class MomentumEquationSystem: public sierra::nalu::EquationSystem
```

Representation of the Momentum conservation equations in 2-D and 3-D.

```
class ContinuityEquationSystem: public sierra::nalu::EquationSystem
```

```
class SpecificDissipationRateEquationSystem: public sierra::nalu::EquationSystem
```

```
class ProjectedNodalGradientEquationSystem: public sierra::nalu::EquationSystem
```

Public Functions

```
void solve_and_update ()
```

Assemble the LHS and RHS and perform linear solve for prescribed number of iterations.

This method is invoked in *EquationSystems::solve_and_update* method as shown below

```
pre_iter_work();
// Iterate over all equation systems
for (auto eqsys: equationSystems_) {
    eqsys->pre_iter_work();
    eqsys->solve_and_update();           //<<<< Assemble and solve system
    eqsys->post_iter_work();
}
```

(continues on next page)

(continued from previous page)

```

}
post_iter_work();

```

See [EquationSystems::solve_and_update](#)

class EquationSystems

A collection of Equations to be solved on a *Realm*.

EquationSystems holds a vector of *EquationSystem* instances representing the equations that are being solved in a given *Realm* and is responsible for the management of the solve and update of the various field quantities in a given timestep.

See [EquationSystems::solve_and_update](#)

Public Functions

bool **solve_and_update** ()

Solve and update the state of all variables for a given timestep.

This method is responsible for executing setup actions before calling solve, performing the actual solve, updating the solution, and performing post-solve actions after the solution has been updated. To provide sufficient granularity and control of this pre- and post- solve actions, the solve method uses the following series of steps:

```

// Perform tasks for this timestep before any Equation system is called
pre_iter_work();
// Iterate over all equation systems
for (auto eqsys: equationSystems_) {
    eqsys->pre_iter_work();
    eqsys->solve_and_update();
    eqsys->post_iter_work();
}
// Perform tasks after all equation systems have updated
post_iter_work();

```

Tasks that require to be performed before any equation system is solved for needs to be registered to `preIterAlgDriver_` on *EquationSystems*, similarly for post-solve tasks. And actions to be performed immediately before individual equation system solves need to be registered in *EquationSystem::preIterAlgDriver_*.

See [pre_iter_work\(\)](#), [post_iter_work\(\)](#), [EquationSystem::pre_iter_work\(\)](#),

See [EquationSystem::post_iter_work\(\)](#)

void **pre_iter_work** ()

Perform necessary setup tasks that affect all *EquationSystem* instances at a given timestep.

See [EquationSystems::solve_and_update\(\)](#)

void **post_iter_work** ()

Perform necessary actions once all *EquationSystem* instances have been updated for the prescribed number of *outer iterations* at a given timestep.

See [EquationSystems::solve_and_update\(\)](#)

Public Members

`std::vector<AlgorithmDriver*> preIterAlgDriver_`

A list of tasks to be performed before all *EquationSystem::solve_and_update*.

`std::vector<AlgorithmDriver*> postIterAlgDriver_`

A list of tasks to be performed after all *EquationSystem::solve_and_update*.

`int numOversetItersDefault_ = {1}`

Default number of overset coupling iterations.

This parameter controls the global settings for *decoupled overset* simulations. User can override this for individual equations by specifying the values for the specific equation system.

`bool decoupledOversetGlobalFlag_ = {false}`

Global flag indicating whether decoupled overset is used for all equation systems in this *Realm*.

User can override this for individual equation systems by using the appropriate input options.

2.2.3 Linear Solvers and Systems Interface

Linear Systems

class LinearSystem

Subclassed by *sierra::nalu::HyprLinearSystem*, *sierra::nalu::TpetraLinearSystem*,
sierra::nalu::TpetraSegregatedLinearSystem

Public Functions

virtual void buildDirichletNodeGraph (**const** stk::mesh::PartVector&)

Process nodes that belong to Dirichlet-type BC.

virtual void buildDirichletNodeGraph (**const** std::vector<stk::mesh::Entity>&)

Process nodes as belonging to a Dirichlet-type row.

See the documentation/implementation of *sierra::nalu::FixPressureAtNodeAlgorithm* for an example of this use case.

virtual void resetRows (**const** std::vector<stk::mesh::Entity> &*nodeList*, **const** unsigned *beginPos*, **const** unsigned *endPos*, **const** double *diag_value* = 0.0, **const** double *rhs_residual* = 0.0) = 0

Reset LHS and RHS for the given set of nodes to 0.

Parameters

- *nodeList*: A list of STK node entities whose rows are zeroed out
- *beginPos*: Starting index (usually 0)
- *endPos*: Terminating index (1 for scalar quantities; nDim for vectors)

class TpetraLinearSystem : **public** *sierra::nalu::LinearSystem*

Public Functions

virtual void resetRows (**const** std::vector<stk::mesh::Entity> &nodeList, **const** unsigned beginPos, **const** unsigned endPos, **const** double diag_value = 0.0, **const** double rhs_residual = 0.0)
Reset LHS and RHS for the given set of nodes to 0.

Parameters

- nodeList: A list of STK node entities whose rows are zeroed out
- beginPos: Starting index (usually 0)
- endPos: Terminating index (1 for scalar quantities; nDim for vectors)

class HypreLinearSystem: public sierra::nalu::LinearSystem

Nalu interface to populate a Hypre Linear System.

This class provides an interface to the HYPRE IJMatrix and IJVector data structures. It is responsible for creating, resetting, and destroying the Hypre data structures and provides the *HypreLinearSystem::sumInto* interface used by Nalu Kernels and SupplementalAlgorithms to populate entries into the linear system. The *HypreLinearSystem::solve* method interfaces with *sierra::nalu::HypreDirectSolver* that is responsible for the actual solution of the system using the required solver and preconditioner combination.

Subclassed by sierra::nalu::HypreUVWLinearSystem

Public Functions

HypreLinearSystem (*Realm* &realm, **const** unsigned numDof, *EquationSystem* *eqSys, *LinearSolver* *linearSolver)

Parameters

- [in] realm: The realm instance that holds the *EquationSystem* being solved
- [in] numDof: The degrees of freedom for the equation system created (Default: 1)
- [in] eqSys: The equation system instance
- [in] linearSolver: Handle to the *HypreDirectSolver* instance

virtual void buildDirichletNodeGraph (**const** stk::mesh::PartVector&)
Tag rows that must be handled as a Dirichlet BC node.

Parameters

- [in] partVec: List of parts that contain the Dirichlet nodes

virtual void buildDirichletNodeGraph (**const** std::vector<stk::mesh::Entity>&)
Tag rows that must be handled as a Dirichlet node.

See sierra::nalu::FixPressureAtNodeAlgorithm

Parameters

- [in] entities: List of nodes where Dirichlet conditions are applied

virtual void loadComplete ()

Finalize construction of the linear system matrix and rhs vector.

This method calls the appropriate Hypr functions to assemble the matrix and rhs in a parallel run, as well as registers the matrix and rhs with the solver preconditioner.

virtual void zeroSystem ()

Reset the matrix and rhs data structures for the next iteration/timestep.

virtual int solve (stk::mesh::FieldBase **linearSolutionField*)

Solve the system $Ax = b$.

The solution vector is returned in linearSolutionField

Parameters

- [out] linearSolutionField: STK field where the solution is populated

double **copy_hypre_to_stk** (stk::mesh::FieldBase *)

Helper method to transfer the solution from a HYPRE_IJVector instance to the STK field data instance.

virtual void applyDirichletBCs (stk::mesh::FieldBase **solutionField*, stk::mesh::FieldBase **bcValuesField*, **const** stk::mesh::PartVector &*parts*, **const** unsigned *beginPos*, **const** unsigned *endPos*)

Populate the LHS and RHS for the Dirichlet rows in linear system.

virtual void sumInto (**const** std::vector<stk::mesh::Entity> &*sym_meshobj*, std::vector<int> &*scratchIds*, std::vector<double> &*scratchVals*, **const** std::vector<double> &*rhs*, **const** std::vector<double> &*lhs*, **const** char **trace_tag*)

Update coefficients of a particular row(s) in the linear system.

The core method of this class, it updates the matrix and RHS based on the inputs from the various algorithms. Note that, unlike *TpetraLinearSystem*, this method skips over the fringe points of Overset mesh and the Dirichlet nodes rather than resetting them afterward.

This overloaded method deals with classic SupplementalAlgorithms

Parameters

- [in] sym_meshobj: A list of STK node entities
- [in] scratchIds: Work array for row IDs
- [in] scratchVals: Work array for row entries
- [in] rhs: Array containing RHS entries to be summed into [numEntities * numDof]
- [in] lhs: Array containing LHS entries to be summed into. [numEntities * numDof * numEntities * numDof]
- [in] trace_tag: Debugging message

virtual void resetRows (**const** std::vector<stk::mesh::Entity> &*nodeList*, **const** unsigned *beginPos*, **const** unsigned *endPos*, **const** double *diag_value*, **const** double *rhs_residual*)

Reset LHS and RHS for the given set of nodes to 0.

Parameters

- nodeList: A list of STK node entities whose rows are zeroed out
- beginPos: Starting index (usually 0)

- `endPos`: Terminating index (1 for scalar quantities; `nDim` for vectors)

```
class HypreLinSysCoeffApplier : public sierra::nalu::CoeffApplier
```

Public Members

```
stk::mesh::NgpMesh ngpMesh_  
    mesh  
  
NGPHyreIDFieldType ngpHypreGlobalId_  
    stk mesh field for the Hypre Global Id  
  
unsigned numDof_ = 0  
    number of degrees of freedom  
  
unsigned nDim_ = 0  
    number of rhs vectors  
  
HypreIntType iLower_ = 0  
    The lowest row owned by this MPI rank.  
  
HypreIntType iUpper_ = 0  
    The highest row owned by this MPI rank.  
  
HypreIntType num_rows_owned_  
    Data structures for the owned CSR Matrix and RHS Vector(s)  
  
HypreIntType num_rows_shared_  
    Data structures for the shared CSR Matrix and RHS Vector(s)  
  
PeriodicNodeMap periodic_node_to_hyre_id_  
    Auxilliary Data structures.  
    map of the periodic nodes to hypre ids  
  
HypreIntTypeViewScalar checkSkippedRows_  
    Flag indicating that sumInto should check to see if rows must be skipped.  
  
HypreIntTypeUnorderedMap skippedRowsMap_  
    unordered map for skipped rows  
  
HypreIntTypeUnorderedMap oversetRowsMap_  
    unordered map for overset rows  
  
HypreLinSysCoeffApplier *devicePointer_  
    this is the pointer to the device function ... that assembles the lists  
  
HypreIntType num_mat_overset_pts_owned_  
    number of points in the overset data structures
```

Linear Solvers Interface

```
class LinearSolver
```

An abstract representation of a linear solver in Nalu.

Defines the basic API supported by the linear solvers for use within Nalu. See concrete implementations such as *sierra::nalu::TpetraLinearSolver* for more details.

Subclassed by *sierra::nalu::HypreDirectSolver*, *sierra::nalu::TpetraLinearSolver*

Public Functions

virtual PetraType **getType** () = 0
 Type of solver instance as defined in `sierra::nalu::PetraType`.

virtual void **destroyLinearSolver** () = 0
 Utility method to cleanup solvers during simulation.

bool &**recomputePreconditioner** ()
 Flag indicating whether the preconditioner is recomputed on each invocation.

bool &**reusePreconditioner** ()
 Flag indicating whether the preconditioner is reused on each invocation.

void **zero_timer_precond** ()
 Reset the preconditioner timer to 0.0 for future accumulation.

double **get_timer_precond** ()
 Get the preconditioner timer for the last invocation.

bool &**activeMueLu** ()
 Flag indicating whether the user has activated MueLU.

LinearSolverConfig ***getConfig** ()
 Get the solver configuration specified in the input file.

Public Members

std::string **name_**
 User-friendly identifier for this particular solver instance.

class **TpetraLinearSolver** : **public** `sierra::nalu::LinearSolver`

Public Functions

TpetraLinearSolver (std::string *solverName*, *TpetraLinearSolverConfig* **config*, **const** Teuchos::RCP<Teuchos::ParameterList> *params*, **const** Teuchos::RCP<Teuchos::ParameterList> *paramsPrecond*, *LinearSolvers* **linearSolvers*)

Parameters

- [in] *solverName*: The name of the solver
- [in] *config*: Solver configuration

virtual void **destroyLinearSolver** ()
 Utility method to cleanup solvers during simulation.

void **setMueLu** ()
 Initialize the MueLU preconditioner before solve.

int **residual_norm** (int *whichNorm*, Teuchos::RCP<LinSys::MultiVector> *sln*, double &*norm*)
 Compute the norm of the non-linear solution vector.

Parameters

- [in] `whichNorm`: [0, 1, 2] norm to be computed
- [in] `sln`: The solution vector
- [out] `norm`: The norm of the solution vector

int **solve** (Teuchos::RCP<LinSys::MultiVector> *sln*, int &*iterationCount*, double &*scaledResidual*, bool *isFinalOuterIter*)
Solve the linear system $Ax = b$.

Parameters

- [out] `sln`: The solution vector
- [out] `iterationCount`: The number of linear solver iterations to convergence
- [out] `scaledResidual`: The final residual norm
- [in] `isFinalOuterIter`: Is this the final outer iteration

virtual PetraType **getType** ()

Type of solver instance as defined in `sierra::nalu::PetraType`.

class **HypreDirectSolver** : **public** `sierra::nalu::LinearSolver`

Nalu interface to Hypre Solvers and Preconditioners.

This class is responsible creation, initialization, execution, and clean up of Hypre solver and preconditioner data structures during the simulation. It provides an abstraction layer so that the user can choose different Hypre solvers via input parameters. This class interacts with rest of Nalu solely via `sierra::nalu::HypreLinearSystem`. The configuration of Hypre solver is controlled via user input parameters processed in `sierra::nalu::HypreLinearSolverConfig`

Users are referred to the [Hypre Reference Manual](#) for detailed documentation on the Hypre functions and data structures used in this class.

Subclassed by `sierra::nalu::HypreUVWSolver`

Public Functions

virtual void **destroyLinearSolver** ()

Clean up Hypre data structures during simulation.

int **solve** (int&, double&, bool)

Solves the linear system and updates the solution vector.

Parameters

- `iters`: The number of linear iterations performed
- `norm`: The norm of the final relative residual

virtual PetraType **getType** ()

Return the type of solver instance.

virtual void **set_initialize_solver_flag** ()

public API for resetting the flag for how often the preconditioner is recomputed

Public Members

HYPRE_ParCSRMatrix **parMat_**
Instance of the Hype parallel matrix.

HYPRE_ParVector **parRhs_**
Instance of the Hype parallel RHS vector.

HYPRE_ParVector **parSln_**
Instance of Hype parallel solution vector.

class LinearSolvers

Collection of solvers and their associated configuration.

This class performs the following actions within a Nalu simulation:

- Parse the `linear_solvers` section and create a mapping of user-defined configurations.
- Create solvers for specific equation system and update the mapping

Public Functions

void **load**(const YAML::Node &node)
Parse the `linear_solvers` section from Nalu input file.

LinearSolver ***create_solver**(std::string solverBlockName, const std::string realmName, EquationType theEQ)
Create a solver for the *EquationSystem*.

Parameters

- [in] solverBlockName: The name specified in the input file, e.g., solve_scalar
- [in] theEQ: The type of equation

Public Members

SolverMap **solvers_**
Mapping of solver instances to the EquationType.

SolverTpetraConfigMap **solverTpetraConfig_**
A lookup table of solver configurations against the names provided in the input file when the type is tpetra

HypreSolverConfigMap **solverHypreConfig_**
A lookup table of solver configurations against the names provided in the input file when type is hypre or tpetra_hypre

Simulation &**sim_**
Reference to the *sierra::nalu::Simulation* instance.

Solver Configuration

class LinearSolverConfig

Subclassed by *sierra::nalu::HypreLinearSolverConfig*, *sierra::nalu::TpetraLinearSolverConfig*

Public Functions

bool **reuseLinSysIfPossible** () const

User flag indicating whether equation systems must attempt to reuse linear system data structures even for cases with mesh motion.

This option only affects decoupled overset system solves where the matrix graph doesn't change, only the entries within the graph. This can be controlled on a per-solver basis.

class TpetraLinearSolverConfig : public sierra::nalu::LinearSolverConfig

class HypreLinearSolverConfig : public sierra::nalu::LinearSolverConfig

User configuration parameters for Hypre solvers and preconditioners.

Public Functions

virtual void **load** (const YAML::Node&)

Process and validate the user inputs and register calls to appropriate Hypre functions to configure the solver and preconditioner.

2.2.4 CVFEM and FEM Interface

class MasterElement

Subclassed by sierra::nalu::Edge2DSCS, sierra::nalu::Edge32DSCS, sierra::nalu::Hex8FEM, sierra::nalu::HexahedralP2Element, sierra::nalu::HexSCS, sierra::nalu::HexSCV, sierra::nalu::PyrSCS, sierra::nalu::PyrSCV, sierra::nalu::Quad3DSCS, sierra::nalu::Quad42DSCS, sierra::nalu::Quad42DSCV, sierra::nalu::Quad93DSCS, sierra::nalu::QuadrilateralP2Element, sierra::nalu::TetSCS, sierra::nalu::TetSCV, sierra::nalu::Tri32DSCS, sierra::nalu::Tri32DSCV, sierra::nalu::Tri3DSCS, sierra::nalu::WedSCS, sierra::nalu::WedSCV

3-D Topologies

class HexSCV : public sierra::nalu::MasterElement

class HexSCS : public sierra::nalu::MasterElement

class TetSCV : public sierra::nalu::MasterElement

class TetSCS : public sierra::nalu::MasterElement

class PyrSCV : public sierra::nalu::MasterElement

class PyrSCS : public sierra::nalu::MasterElement

class WedSCV : public sierra::nalu::MasterElement

class WedSCS : public sierra::nalu::MasterElement

class Hex27SCV : public sierra::nalu::HexahedralP2Element

class Hex27SCS : public sierra::nalu::HexahedralP2Element

class Hex8FEM : public sierra::nalu::MasterElement

class Quad3DSCS : public sierra::nalu::MasterElement

class Quad93DSCS : public sierra::nalu::MasterElement

class Tri3DSCS : public sierra::nalu::MasterElement

2-D Topologies

```
class Quad42DSCV : public sierra::nalu::MasterElement
class Quad42DSCS : public sierra::nalu::MasterElement
class Tri32DSCV : public sierra::nalu::MasterElement
class Tri32DSCS : public sierra::nalu::MasterElement
```

Higher-order Element Topologies

Warning: doxygenclass: Cannot find class “sierra::nalu::HigherOrderHexSCV” in doxygen xml output for project “nalu” from directory: ./doxygen/xml

Warning: doxygenclass: Cannot find class “sierra::nalu::HigherOrderHexSCS” in doxygen xml output for project “nalu” from directory: ./doxygen/xml

Warning: doxygenclass: Cannot find class “sierra::nalu::HigherOrderQuad2DSCV” in doxygen xml output for project “nalu” from directory: ./doxygen/xml

Warning: doxygenclass: Cannot find class “sierra::nalu::HigherOrderQuad2DSCS” in doxygen xml output for project “nalu” from directory: ./doxygen/xml

2.2.5 Actuator Sources

The *sierra::nalu::ActuatorLineFAST* class is a child class of the generic *sierra::nalu::Actuator* class that couples Nalu with OpenFAST for actuator line simulations of wind turbines.

```
class Actuator
    Subclassed by sierra::nalu::ActuatorFAST, sierra::nalu::ActuatorSimple
class ActuatorLineFAST : public sierra::nalu::ActuatorFAST
```

2.2.6 Auxiliary Functions

```
class AuxFunction
    Subclassed by sierra::nalu::BoundaryLayerPerturbationAuxFunction, sierra::nalu::BoussinesqNonIsoTemperatureAuxFunction,
    sierra::nalu::BoussinesqNonIsoVelocityAuxFunction, sierra::nalu::CappingInversionTemperatureAuxFunction,
    sierra::nalu::ConstantAuxFunction, sierra::nalu::ConvectingTaylorVortexPressureAuxFunction,
    sierra::nalu::ConvectingTaylorVortexPressureGradAuxFunction, sierra::nalu::ConvectingTaylorVortexVelocityAuxFunction,
    sierra::nalu::FlowPastCylinderTempAuxFunction, sierra::nalu::GaussJetVelocityAuxFunction,
    sierra::nalu::KovasznyPressureAuxFunction, sierra::nalu::KovasznyPressureGradientAuxFunction,
    sierra::nalu::KovasznyVelocityAuxFunction, sierra::nalu::LinearRampMeshDisplacementAuxFunction,
    sierra::nalu::OneTwoTenVelocityAuxFunction, sierra::nalu::PerturbedShearLayerMixFracAuxFunction,
    sierra::nalu::PerturbedShearLayerVelocityAuxFunction, sierra::nalu::RayleighTaylorMixFracAuxFunction,
    sierra::nalu::SinMeshDisplacementAuxFunction, sierra::nalu::SinProfileChannelFlowVelocityAuxFunction,
```

sierra::nalu::SteadyTaylorVortexGradPressureAuxFunction, sierra::nalu::SteadyTaylorVortexPressureAuxFunction,
sierra::nalu::SteadyTaylorVortexVelocityAuxFunction, sierra::nalu::SteadyThermal3dContactAuxFunction,
sierra::nalu::SteadyThermal3dContactDtDxAuxFunction, sierra::nalu::SteadyThermalContactAuxFunction,
sierra::nalu::TaylorGreenPressureAuxFunction, sierra::nalu::TaylorGreenVelocityAuxFunction,
sierra::nalu::TornadoAuxFunction, sierra::nalu::VariableDensityMixFracAuxFunction,
sierra::nalu::VariableDensityNonIsoTemperatureAuxFunction, sierra::nalu::VariableDensityPressureAuxFunction,
sierra::nalu::VariableDensityVelocityAuxFunction, sierra::nalu::WindEnergyPowerLawAuxFunction,
sierra::nalu::WindEnergyTaylorVortexAuxFunction, sierra::nalu::WindEnergyTaylorVortexPressureAuxFunction,
sierra::nalu::WindEnergyTaylorVortexPressureGradAuxFunction

ABL Utilities

class BoundaryLayerPerturbationAuxFunction : public sierra::nalu::AuxFunction
Add sinusoidal perturbations to the velocity field.

This function is used as an initial condition, primarily in Atmospheric Boundary Layer (ABL) flows, to trigger transition to turbulent flow during ABL precursor simulations.

Steady Taylor Vortex

class SteadyTaylorVortexVelocityAuxFunction : public sierra::nalu::AuxFunction
class SteadyTaylorVortexPressureAuxFunction : public sierra::nalu::AuxFunction
class SteadyTaylorVortexGradPressureAuxFunction : public sierra::nalu::AuxFunction
class SteadyTaylorVortexMomentumSrcElemSuppAlg : public sierra::nalu::SupplementalAlgorithm
class SteadyTaylorVortexMomentumSrcNodeSuppAlg : public sierra::nalu::SupplementalAlgorithm

Convecting Taylor Vortex

class ConvectingTaylorVortexVelocityAuxFunction : public sierra::nalu::AuxFunction
class ConvectingTaylorVortexPressureAuxFunction : public sierra::nalu::AuxFunction
class ConvectingTaylorVortexPressureGradAuxFunction : public sierra::nalu::AuxFunction

Kovasznay 2-D Flow

class KovasznayVelocityAuxFunction : public sierra::nalu::AuxFunction
class KovasznayPressureAuxFunction : public sierra::nalu::AuxFunction
class KovasznayPressureGradientAuxFunction : public sierra::nalu::AuxFunction

Steady Thermal MMS (2-D and 3-D)

class SteadyThermal3dContactAuxFunction : public sierra::nalu::AuxFunction
class SteadyThermal3dContactDtDxAuxFunction : public sierra::nalu::AuxFunction
template<typename AlgTraits>
class SteadyThermal3dContactSrcElemKernel : public sierra::nalu::Kernel

Public Functions

```
virtual void execute (SharedMemView<DoubleType **>&, SharedMemView<DoubleType *>&,
    ScratchViews<DoubleType>&)
```

Execute the kernel within a Kokkos loop and populate the LHS and RHS for the linear solve.

```
class SteadyThermal3dContactSrcElemSuppAlgDep : public sierra::nalu::SupplementalAlgorithm
```

```
class SteadyThermalContact3DSrcNodeSuppAlg : public sierra::nalu::SupplementalAlgorithm
```

```
class SteadyThermalContactAuxFunction : public sierra::nalu::AuxFunction
```

```
class SteadyThermalContactSrcElemSuppAlg : public sierra::nalu::SupplementalAlgorithm
```

```
class SteadyThermalContactSrcNodeSuppAlg : public sierra::nalu::SupplementalAlgorithm
```

Mesh Motion/Displacement Utilities

```
class LinearRampMeshDisplacementAuxFunction : public sierra::nalu::AuxFunction
```

```
class SinMeshDisplacementAuxFunction : public sierra::nalu::AuxFunction
```

Warning: doxygenclass: Cannot find class “sierra::nalu::WindEnergyAuxFunction” in doxygen xml output for project “nalu” from directory: ./doxygen/xml

2.2.7 Post-Processing Utilities

```
class TurbulenceAveragingPostProcessing
```

Post-processing to collect various types of statistics on flow fields.

This class implements Reynolds and Favre averaging as well as other useful quantities relevant to analyzing turbulent flows.

Currently supported:

- Reynolds and Favre averaging of flow variables
- TKE and stress computation
- Vorticity, Q-criterion, lambda-ci calculation

Public Types

```
enum AveragingType
```

Type of time filter averaging applied.

Values:

```
NALU_CLASSIC = 0
```

Classic Nalu implementation (saw-tooth reset)

```
MOVING_EXPONENTIAL
```

Moving exponential window averaging.

```
class DataProbePostProcessing
```

```
class SolutionNormPostProcessing
```

```
class SurfaceForceAndMomentAlgorithm : public sierra::nalu::Algorithm
```

```
class SurfaceForceAndMomentWallFunctionAlgorithm: public sierra::nalu::Algorithm
```

2.3 Writing Developer Documentation

Developer documentation should be written using Doxygen annotations directly in the source code. This allows the documentation to live with the code essentially as comments that Doxygen is able to extract automatically into a more human readable form. Doxygen requires special syntax markers to indicate comments that should be processed as inline documentation vs. generic comments in the source code. The [Doxygen manual](#) provides detailed information on the various markers available to tailor the formatting of auto-generated documentation. It is recommended that users document the classes and methods in the header file. A sample header file with specially formatted comments is shown below. You can download a copy of the file.

Listing 2.1: Sample C++ header file showing inline documentation using specially formatted comments.

```
/** @file example.h
 *  @brief Brief description of a documented file.
 *
 *  Longer description of a documented file.
 */

/** Here is a brief description of the example class.
 *
 *  This is a more in-depth description of the class.
 *  This class is meant as an example.
 *  It is not useful by itself, rather its usefulness is only a
 *  function of how much it helps the reader. It is in a sense
 *  defined by the person who reads it and otherwise does
 *  not exist in any real form.
 *
 *  @note This is a note.
 */

#ifndef EXAMPLECLASS_H
#define EXAMPLECLASS_H

class ExampleClass
{
public:

    /// Create an ExampleClass.
    ExampleClass();

    /** Create an ExampleClass with lot's of intial values.
     *
     *  @param a This is a description of parameter a.
     *  @param b This is a description of parameter b.
     *
     *  The distance between  $(x_1, y_1)$  and  $(x_2, y_2)$  is
     *   $\sqrt{(x_2 - x_1)^2 + (y_2 - y_1)^2}$ .
     */
    ExampleClass(int a, float b);
```

(continues on next page)

(continued from previous page)

```

/** ExampleClass destructor description.
 */
~ExampleClass();

/// This method does something.
void DoSomething();

/**
 * This is a method that does so
 * much that I must write an epic
 * novel just to describe how much
 * it truly does.
 */
void DoNothing();

/** Brief description of a useful method.
 * @param level An integer setting how useful to be.
 * @return Description of the output.
 *
 * This method does unbelievably useful things.
 * And returns exceptionally useful results.
 * Use it everyday with good health.
 * \f[
 * |I_2|=\left| \int_0^T \psi(t)
 * \left\{
 * u(a,t)-
 * \int_{\gamma(t)}^a
 * \frac{d\theta}{k(\theta,t)}
 * \int_a^\theta c(\xi)u_t(\xi,t)\,d\xi
 * \right\} dt
 * \right|
 * \f]
 */
void* VeryUsefulMethod(bool level);

/** Brief description of a useful method.
 * @param level An integer setting how useful to be.
 * @return Description of the output.
 *
 * - Item 1
 *
 * More text for this item.
 *
 * - Item 2
 *   + nested list item.
 *   + another nested item.
 * - Item 3
 *
 * # Markdown Example
 * [Here is a link.](http://www.google.com/)
 */
void* AnotherMethod(bool level);

protected:
/** The protected methods can be documented and extracted too.
 *
 */

```

(continues on next page)

(continued from previous page)

```

    void SomeProtectedMethod();

private:

    const char* fQuestion; ///< The question
    int fAnswer;           ///< The answer

}; // End of class ExampleClass

#endif // EXAMPLE_H

```

Once processed by Doxygen and embedded in Sphinx, the resulting documentation of the class looks as shown below:

class ExampleClass

Here is a brief description of the example class.

This is a more in-depth description of the class. This class is meant as an example. It is not useful by itself, rather its usefulness is only a function of how much it helps the reader. It is in a sense defined by the person who reads it and otherwise does not exist in any real form.

Note This is a note.

Public Functions

ExampleClass()

Create an *ExampleClass*.

ExampleClass(int a, float b)

Create an *ExampleClass* with lot's of initial values.

The distance between (x_1, y_1) and (x_2, y_2) is $\sqrt{(x_2 - x_1)^2 + (y_2 - y_1)^2}$.

Parameters

- a: This is a description of parameter a.
- b: This is a description of parameter b.

~ExampleClass()

ExampleClass destructor description.

void DoSomething()

This method does something.

void DoNothing()

This is a method that does so much that I must write an epic novel just to describe how much it truly does.

void *VeryUsefulMethod(bool level)

Brief description of a useful method.

This method does unbelievably useful things. And returns exceptionally useful results. Use it everyday with good health.

$$|I_2| = \left| \int_0^T \psi(t) \left\{ u(a, t) - \int_{\gamma(t)}^a \frac{d\theta}{k(\theta, t)} \int_a^\theta c(\xi) u_t(\xi, t) d\xi \right\} dt \right|$$

Return Description of the output.

Parameters

- `level`: An integer setting how useful to be.

void ***AnotherMethod** (bool *level*)

Brief description of a useful method.

- Item 1
More text for this item.
- Item 2
 - nested list item.
 - another nested item.
- Item 3

Return Description of the output.

Parameters

- `level`: An integer setting how useful to be.

Markdown Example

Here is a [link](#).

Protected Functions

void **SomeProtectedMethod** ()

The protected methods can be documented and extracted too.

Private Members

const char ***fQuestion**

The question.

int **fAnswer**

The answer.

2.4 Writing User Documentation

This documentation is written in Sphinx and is generated automatically on the <http://nalu-wind.readthedocs.io> website every time the [Nalu-Wind Github repo](#) is updated. This documentation can also be built locally on your machine by using the instructions [here](#). Sphinx uses restructured text (RST) to generate the documentation in many other formats, such as this html version. Refer to the primer on writing restructured text [here](#).

2.5 Building the Documentation

This document describes how to build Nalu-Wind's documentation. The documentation is based on the use of Doxygen, Sphinx, and Doxylink. Therefore we will need to install these tools as well as some extensions of Sphinx that are utilized.

2.5.1 Install the Tools

Install CMake, Doxygen, Sphinx, Doxylink, and the extensions used. Doxygen uses the `dot` application installed with GraphViz. Sphinx uses a combination of extensions installed with `pip install` as well as some that come with Nalu-Wind located in the `_extensions` directory. Using Homebrew on Mac OS X, this would look something like:

```
brew install cmake
brew install python
brew install doxygen
brew install graphviz
pip2 install sphinx
pip2 install sphinxcontrib-bibtex
pip2 install breathe
pip2 install sphinx_rtd_theme
```

On Linux, CMake, Python, Doxygen, and GraphViz could be installed using your package manager, e.g. `sudo apt-get install cmake`.

2.5.2 Run CMake Configure

In the [Nalu-Wind repository](#) checkout, create your own or use the `build` directory that already exists in the repo. Change to your designated build directory and run CMake with `-DENABLE_DOCUMENTATION` on. For example:

```
cmake -DTrilinos_DIR:PATH=$(spack location -i nalu-trilinos) \
      -DYAML_DIR:PATH=$(spack location -i yaml-cpp) \
      -DCMAKE_BUILD_TYPE=RELEASE \
      -DENABLE_DOCUMENTATION:BOOL=ON \
      ..
```

If all of the main tools are found successfully, CMake should configure with the ability to build the documentation. If Sphinx or Doxygen aren't found, the configure will skip the documentation.

2.5.3 Make the Docs

In your designated build directory, issue the command `make docs` which should first build the Doxygen documentation and then the Sphinx documentation. If this completes successfully, the entry point to the documentation should be in `build/docs/html/index.html`.

2.6 Developer Workflow

This document describes a suggested developer workflow for Nalu-Wind.

2.7 Nalu Style Guide

1. No tabs. Remove them from your editor. Better yet, use eclipse and follow the xml style. Use the format [here](#).
2. Use underscores for private data, e.g., `const double thePrivateData_`.
3. Use camel case for data members and classes unless it is silly (you get the idea).
4. Camel case on Class names always; non camel case for methods, e.g.,


```
const double Realm::get_me() {  
    return hereIAm_; // hmmm... silly? your call  
}
```

5. Use `const` when possible, however, do not try to be a member of the ‘const’ police force.
6. We need logic to launch some special physics. Try to avoid run time logic by designing with polymorphic/templates.
7. When possible, add classes that manage loading, field registration, setup and execute, e.g., `SolutionNormPostProcessing`, etc.

2.8 Contributing to Nalu-Wind

1. There is no rush to push. We only support production tested capability. Better yet, perform code verification and unit testing.
2. Always run the full regression test suite. No exceptions.
3. Peer review when fully appropriate (ask for a pull request).
4. If adding a new feature, include a regression test for this feature. Refer to the section of this documentation on adding a test [here](#).

NALU-WIND - THEORY MANUAL

Nalu-Wind represents a generalized unstructured, massively parallel, variable density turbulent flow capability designed for energy applications. This code base began as an effort to prototype Sierra Toolkit, [EWS+10], usage along with direct parallel matrix assembly to the Trilinos, [HBH+03], Epetra and Tpetra data structure. However, the simulation tool has evolved as a tool to support a variety of research projects germane to the energy sector including wind aerodynamic prediction and traditional gas-phase combustion applications.

3.1 Low Mach Number Derivation

The low Mach number equations are a subset of the fully compressible equations of motion (momentum, continuity and energy), admitting large variations in gas density while remaining acoustically incompressible. The low Mach number equations are preferred over the full compressible equations for low speed flow problems as the acoustics are of little consequence to the overall simulation accuracy. The technique avoids the need to resolve fast-moving acoustic signals. Derivations of the low Mach number equations can be found in Rehman and Baum, [RB78], or Paolucci, [Pao82].

The equations are derived from the compressible equations using a perturbation expansion in terms of the lower limit of the Mach number squared; hence the name. The asymptotic expansion leads to a splitting of pressure into a spatially constant thermodynamic pressure and a locally varying dynamic pressure. The dynamic pressure is decoupled from the thermodynamic state and cannot propagate acoustic waves. The thermodynamic pressure is used in the equation of state and to determine thermophysical properties. The thermodynamic pressure can vary in time and can be calculated using a global energy balance.

3.1.1 Asymptotic Expansion

The asymptotic expansion for the low Mach number equations begins with the full compressible equations in Cartesian coordinates. The equations are the minimum set required to propagate acoustic waves. The equations are written in divergence form using Einstein notation (summation over repeated indices):

$$\begin{aligned}\frac{\partial \rho}{\partial t} + \frac{\partial \rho u_j}{\partial x_j} &= 0, \\ \frac{\partial \rho u_i}{\partial t} + \frac{\partial \rho u_j u_i}{\partial x_j} + \frac{\partial P}{\partial x_i} &= \frac{\partial \tau_{ij}}{\partial x_j} + \rho g_i, \\ \frac{\partial \rho E}{\partial t} + \frac{\partial \rho u_j H}{\partial x_j} &= -\frac{\partial q_j}{\partial x_j} + \frac{\partial u_i \tau_{ij}}{\partial x_j} + \rho u_i g_i.\end{aligned}$$

The primitive variables are the velocity components, u_i , the pressure, P , and the temperature T . The viscous shear stress tensor is τ_{ij} , the heat conduction is q_i , the total enthalpy is H , the total internal energy is E , the density is ρ ,

and the gravity vector is g_i . The total internal energy and total enthalpy contain the kinetic energy contributions. The equations are closed using the following models and definitions:

$$\begin{aligned} P &= \rho \frac{R}{W} T, \\ E &= H - P/\rho, \\ H &= h + \frac{1}{2} u_k u_k, \\ \tau_{ij} &= \mu \left(\frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i} \right) - \frac{2}{3} \mu \frac{\partial u_k}{\partial x_k} \delta_{ij}, \\ q_i &= -k \frac{\partial T}{\partial x_i} \end{aligned}$$

The mean molecular weight of the gas is W , the molecular viscosity is μ , and the thermal conductivity is k . A Newtonian fluid is assumed along with the Stokes hypothesis for the stress tensor.

The equations are scaled so that the variables are all of order one. The velocities, lengths, and times are nondimensionalized by a characteristic velocity, U_∞ , and a length scale, L . The pressure, density, and temperature are nondimensionalized by P_∞ , ρ_∞ , and T_∞ . The enthalpy and energy are nondimensionalized by $C_{p,\infty} T_\infty$. Dimensionless variables are noted by overbars. The dimensionless equations are:

$$\begin{aligned} \frac{\partial \bar{\rho}}{\partial \bar{t}} + \frac{\partial \bar{\rho} \bar{u}_j}{\partial \bar{x}_j} &= 0, \\ \frac{\partial \bar{\rho} \bar{u}_i}{\partial \bar{t}} + \frac{\partial \bar{\rho} \bar{u}_j \bar{u}_i}{\partial \bar{x}_j} + \frac{1}{\gamma \text{Ma}^2} \frac{\partial \bar{P}}{\partial \bar{x}_i} &= \frac{1}{\text{Re}} \frac{\partial \bar{\tau}_{ij}}{\partial \bar{x}_j} + \frac{1}{\text{Fr}_i} \bar{\rho}, \\ \frac{\partial \bar{\rho} \bar{h}}{\partial \bar{t}} + \frac{\partial \bar{\rho} \bar{u}_j \bar{h}}{\partial \bar{x}_j} &= -\frac{1}{\text{Pr}} \frac{1}{\text{Re}} \frac{\partial \bar{q}_j}{\partial \bar{x}_j} + \frac{\gamma - 1}{\gamma} \frac{\partial \bar{P}}{\partial \bar{t}} \\ &\quad + \frac{\gamma - 1}{\gamma} \frac{\text{Ma}^2}{\text{Re}} \frac{\partial \bar{u}_i \bar{\tau}_{ij}}{\partial \bar{x}_j} + \bar{\rho} \bar{u}_i \frac{\gamma - 1}{\gamma} \frac{\text{Ma}^2}{\text{Fr}_i} \\ &\quad - \frac{\gamma - 1}{2} \text{Ma}^2 \left(\frac{\partial \bar{\rho} \bar{u}_k \bar{u}_k}{\partial \bar{t}} + \frac{\partial \bar{\rho} \bar{u}_j \bar{u}_k \bar{u}_k}{\partial \bar{x}_j} \right). \end{aligned}$$

The groupings of characteristic scaling terms are:

$$\begin{aligned} \text{Re} &= \frac{\rho_\infty U_\infty L}{\mu_\infty}, & \text{Reynoldsnnumber}, \\ \text{Pr} &= \frac{C_{p,\infty} \mu_\infty}{k_\infty}, & \text{Prandtlnumber}, \\ \text{Fr}_i &= \frac{u_\infty^2}{g_i L}, & \text{Froudenumber}, \quad g_i \neq 0, \\ \text{Ma} &= \sqrt{\frac{u_\infty^2}{\gamma R T_\infty / W}}, & \text{Machnumber}, \end{aligned}$$

where γ is the ratio of specific heats.

For small Mach numbers, $\text{Ma} \ll 1$, the kinetic energy, viscous work, and gravity work terms can be neglected in the energy equation since those terms are scaled by the square of the Mach number. The inverse of Mach number squared remains in the momentum equations, suggesting singular behavior. In order to explore the singularity, the pressure, velocity and temperature are expanded as asymptotic series in terms of the parameter ϵ :

$$\begin{aligned} \bar{P} &= \bar{P}_0 + \bar{P}_1 \epsilon + \bar{P}_2 \epsilon^2 \dots \\ \bar{u}_i &= \bar{u}_{i,0} + \bar{u}_{i,1} \epsilon + \bar{u}_{i,2} \epsilon^2 \dots \\ \bar{T} &= \bar{T}_0 + \bar{T}_1 \epsilon + \bar{T}_2 \epsilon^2 \dots \end{aligned}$$

The zeroeth-order terms are collected together in each of the equations. The form of the continuity equation stays the same. The gradient of the pressure in the zeroeth-order momentum equations can become singular since it is divided by the characteristic Mach number squared. In order for the zeroeth-order momentum equations to remain well-behaved, the spatial variation of the \bar{P}_0 term must be zero. If the magnitude of the expansion parameter is selected to be proportional to the square of the characteristic Mach number, $\epsilon = \gamma \text{Ma}^2$, then the \bar{P}_1 term can be included in the zeroeth-order momentum equation.

$$\frac{1}{\gamma \text{Ma}^2} \frac{\partial \bar{P}}{\partial x_i} = \frac{\partial}{\partial x_i} \left(\frac{1}{\gamma \text{Ma}^2} \bar{P}_0 + \frac{\epsilon}{\gamma \text{Ma}^2} \bar{P}_1 + \dots \right) = \frac{\partial}{\partial x_i} \left(\bar{P}_1 + \epsilon \bar{P}_2 + \dots \right)$$

The form of the energy equation remains the same, less the kinetic energy, viscous work and gravity work terms. The P_0 term remains in the energy equation as a time derivative. The low Mach number equations are the zeroeth-order equations in the expansion including the P_1 term in the momentum equations. The expansion results in two different types of pressure and they are considered to be split into a thermodynamic component and a dynamic component. The thermodynamic pressure is constant in space, but can change in time. The thermodynamic pressure is used in the equation of state. The dynamic pressure only arises as a gradient term in the momentum equation and acts to enforce continuity. The unsplit dimensional pressure is

$$P = P_{th} + \gamma \text{Ma}^2 P_1,$$

where the dynamic pressure, $p = P - P_{th}$, is related to a pressure coefficient

$$\bar{P}_1 = \frac{P - P_{th}}{\rho_\infty u_\infty^2} P_{th}.$$

The resulting unscaled low Mach number equations are:

$$\begin{aligned} \frac{\partial \rho}{\partial t} + \frac{\partial \rho u_j}{\partial x_j} &= 0, \\ \frac{\partial \rho u_i}{\partial t} + \frac{\partial \rho u_j u_i}{\partial x_j} + \frac{\partial P}{\partial x_i} &= \frac{\partial \tau_{ij}}{\partial x_j} + (\rho - \rho_\infty) g_i, \\ \frac{\partial \rho h}{\partial t} + \frac{\partial \rho u_j h}{\partial x_j} &= -\frac{\partial q_j}{\partial x_j} + \frac{\partial P_{th}}{\partial t}, \end{aligned}$$

where the ideal gas law becomes

$$P_{th} = \rho \frac{R}{W} T.$$

The hydrostatic pressure gradient has been subtracted from the momentum equation, assuming an ambient density of ρ_∞ . The stress tensor and heat conduction remain the same as in the original equations.

3.2 Supported Equation Set

This section provides an overview of the currently supported equation sets. Equations will be described in integral form with assumed Favre averaging. However, the laminar counterparts are supported in the code base and are obtain in the user file by omitting a turbulence model specification.

3.2.1 Conservation of Mass

The continuity equation is always solved in the variable density form.

$$\int \frac{\partial \bar{\rho}}{\partial t} dV + \int \bar{\rho} \tilde{u}_i n_i dS = 0$$

Since Nalu-Wind uses equal-order interpolation (variables are collocated) stabilization is required. The stabilization choice will be developed in the pressure stabilization section.

Note that the use of a low speed compressible formulation requires that the fluid density be computed by an equation of state that uses the thermodynamic pressure. This thermodynamic pressure can either be computed based on a global energy/mass balance or allowed to be spatially varying. By modification of the continuity density time derivative to include the $\frac{\partial \rho}{\partial p}$ sensitivity, an equation that admits acoustic pressure waves is realized.

3.2.2 Conservation of Momentum

The integral form of the Favre-filtered momentum equations used for turbulent transport are

$$\int \frac{\partial \bar{\rho} \tilde{u}_i}{\partial t} dV + \int \bar{\rho} \tilde{u}_i \tilde{u}_j n_j dS = \int \tilde{\sigma}_{ij} n_j dS - \int \tau_{ij}^{sgs} n_j dS + \int (\bar{\rho} - \rho_o) g_i dV + \int f_i dV, \quad (3.1)$$

where the subgrid scale turbulent stress τ_{ij}^{sgs} is defined as

$$\tau_{ij}^{sgs} \equiv \bar{\rho}(\widetilde{u_i u_j} - \tilde{u}_i \tilde{u}_j). \quad (3.2)$$

The term f_i is a body force used to represent additional momentum sources such as wind turbine blades, Coriolis effect, driving forces, etc. The Cauchy stress is provided by,

$$\sigma_{ij} = 2\mu \tilde{S}_{ij}^* - \bar{P} \delta_{ij}$$

and the traceless rate-of-strain tensor defined as follows:

$$\begin{aligned} \tilde{S}_{ij}^* &= \tilde{S}_{ij} - \frac{1}{3} \delta_{ij} \tilde{S}_{kk} \\ &= \tilde{S}_{ij} - \frac{1}{3} \frac{\partial \tilde{u}_k}{\partial x_k} \delta_{ij}. \end{aligned}$$

In a low Mach flow, as described in the low Mach theory section, the above pressure, \bar{P} is the perturbation about the thermodynamic pressure, P^{th} . In a low speed compressible flow, i.e., flows confined to a closed domain with energy or mass addition in which the continuity equation has been modified to accommodate acoustics, this pressure is interpreted at the thermodynamic pressure itself.

For LES, τ_{ij}^{sgs} that appears in Equation (3.1) and defined in Equation (3.2) represents the subgrid stress tensor that must be closed. The deviatoric part of the subgrid stress tensor is defined as

$$\tau_{ij}^{sgs} = \tau_{ij}^{sgs} - \frac{1}{3} \delta_{ij} \tau_{kk}^{sgs} \quad (3.3)$$

where the subgrid turbulent kinetic energy is defined as $\tau_{kk}^{sgs} = 2\bar{\rho}k$. Note that here, k represents the modeled turbulent kinetic energy as is formally defined as,

$$\bar{\rho}k = \frac{1}{2} \bar{\rho}(\widetilde{u_k u_k} - \tilde{u}_k \tilde{u}_k).$$

Model closures can use, Yoshizawa's approach when k is not transported:

$$\tau_{kk}^{sgs} = 2C_I \bar{\rho} \Delta^2 |\tilde{S}|^2.$$

Above, $|\tilde{S}| = \sqrt{2\tilde{S}_{ij}\tilde{S}_{ij}}$.

For low Mach-number flows, a vast majority of the turbulent kinetic energy is contained at resolved scales. For this reason, the subgrid turbulent kinetic energy is not directly treated and, rather, is included in the pressure as an additional normal stress. The Favre-filtered momentum equations then become

$$\begin{aligned} \int \frac{\partial \bar{\rho} \tilde{u}_i}{\partial t} dV + \int \bar{\rho} \tilde{u}_i \tilde{u}_j n_j dS + \int \left(\bar{P} + \frac{2}{3} \bar{\rho} k \right) n_i dS = \\ \int 2(\mu + \mu_t) \left(\tilde{S}_{ij} - \frac{1}{3} \tilde{S}_{kk} \delta_{ij} \right) n_j dS + \int (\bar{\rho} - \rho_o) g_i dV, \end{aligned} \quad (3.4)$$

where LES closure models for the subgrid turbulent eddy viscosity μ_t are either the constant coefficient Smagorinsky, WALE or the constant coefficient k_{sgs} model (see the turbulence section).

Earth Coriolis Force

For simulation of large-scale atmospheric flows, the following Coriolis force term can be added to the right-hand-side of the momentum equation ((3.1)):

$$f_i = -2\bar{\rho}\epsilon_{ijk}\Omega_j u_k. \quad (3.5)$$

Here, Ω is the Earth's angular velocity vector, and ϵ_{ijk} is the Levi-Civita symbol denoting the cross product of the Earth's angular velocity with the local fluid velocity vector. Consider an “East-North-Up” coordinate system on the Earth's surface, with the domain centered on a latitude angle ϕ (changes in latitude within the computational domain are neglected). In this coordinate system, the integrand of (cor-term), or the Coriolis acceleration vector, is

$$2\bar{\rho}\omega \begin{bmatrix} u_n \sin \phi - u_u \cos \phi \\ -u_e \sin \phi \\ u_e \cos \phi \end{bmatrix}, \quad (3.6)$$

where $\omega \equiv ||\Omega||$. Often, in geophysical flows it is assumed that the vertical component of velocity is small and that the vertical component of the acceleration is small relative to gravity, such that the terms containing $\cos \phi$ are neglected. However, there is evidence that this so-called traditional approximation is not valid for some mesoscale atmospheric phenomena cite{Gerkema_etal:08}, and so the full Coriolis term is retained in Nalu-Wind. The implementation proceeds by first finding the velocity vector in the East-North-Up coordinate system, then calculating the Coriolis acceleration vector ((3.6)), then transforming this vector back to the model $x - y - z$ coordinate system. The coordinate transformations are made using user-supplied North and East unit vectors given in the model coordinate system.

Boussinesq Buoyancy Model

In atmospheric and other flows, the density differences in the domain can be small enough as to not significantly affect inertia, but nonetheless the buoyancy term,

$$\int (\bar{\rho} - \rho_o) g_i dV, \quad (3.7)$$

may still be important. The Boussinesq model ignores the effect of density on inertia while retaining the buoyancy term in Equation (3.1). For the purpose of evaluating the buoyant force, the density is approximated as

$$\frac{\rho}{\rho_o} \approx 1 - \beta(T - T_o), \quad (3.8)$$

This leads to a buoyancy body force term that depends on temperature (T), a reference density (ρ_o), a reference temperature (T_o), and a thermal expansion coefficient (β) as

$$\int -\rho_o \beta (T - T_o) g_i dV. \quad (3.9)$$

The flow is otherwise kept as constant density.

3.2.3 Filtered Mixture Fraction

The optional quantity used to identify the chemical state is the mixture fraction, Z . While there are many different definitions of the mixture fraction that have subtle variations that attempt to capture effects like differential diffusion, they can all be interpreted as a local mass fraction of the chemical elements that originated in the fuel stream. The mixture fraction is a conserved scalar that varies between zero in the secondary stream and unity in the primary stream and is transported in laminar flow by the equation,

$$\frac{\partial \rho Z}{\partial t} + \frac{\partial \rho u_j Z}{\partial x_j} = \frac{\partial}{\partial x_j} \left(\rho D \frac{\partial Z}{\partial x_j} \right), \quad (3.10)$$

where D is an effective molecular mass diffusivity.

Applying either temporal Favre filtering for RANS-based treatments or spatial Favre filtering for LES-based treatments yields

$$\int \frac{\partial \bar{\rho} \tilde{Z}}{\partial t} dV + \int \bar{\rho} \tilde{u}_j \tilde{Z} n_j dS = - \int \tau_{Z,j}^{sgs} n_j dS + \int \bar{\rho} D \frac{\partial \tilde{Z}}{\partial x_j} n_j dS, \quad (3.11)$$

where sub-filter correlations have been neglected in the molecular diffusive flux vector and the turbulent diffusive flux vector is defined as

$$\tau_{Z,j}^{sgs} \equiv \bar{\rho} \left(\widetilde{Z u_j} - \tilde{Z} \tilde{u}_j \right).$$

This subgrid scale closure is modeled using the gradient diffusion hypothesis,

$$\tau_{Z,j}^{sgs} = -\bar{\rho} D_t \frac{\partial \tilde{Z}}{\partial x_j},$$

where D_t is the turbulent mass diffusivity, modeled as $\bar{\rho} D_t = \mu_t / Sc_t$ where μ_t is the modeled turbulent viscosity from momentum transport and Sc_t is the turbulent Schmidt number. The molecular mass diffusivity is then expressed similarly as $\bar{\rho} D = \mu / Sc$ so that the final modeled form of the filtered mixture fraction transport equation is

$$\frac{\partial \bar{\rho} \tilde{Z}}{\partial t} + \frac{\partial \bar{\rho} \tilde{u}_j \tilde{Z}}{\partial x_j} = \frac{\partial}{\partial x_j} \left[\left(\frac{\mu}{Sc} + \frac{\mu_t}{Sc_t} \right) \frac{\partial \tilde{Z}}{\partial x_j} \right].$$

In integral form the mixture fraction transport equation is

$$\int \frac{\partial \bar{\rho} \tilde{Z}}{\partial t} dV + \int \bar{\rho} \tilde{u}_j \tilde{Z} n_j dS = \int \left(\frac{\mu}{Sc} + \frac{\mu_t}{Sc_t} \right) \frac{\partial \tilde{Z}}{\partial x_j} n_j dS.$$

3.2.4 Conservation of Energy

The integral form of the Favre-filtered static enthalpy energy equation used for turbulent transport is

$$\begin{aligned} \int \frac{\partial \bar{\rho} \tilde{h}}{\partial t} dV + \int \bar{\rho} \tilde{h} \tilde{u}_j n_j dS = & - \int \bar{q}_j n_j dS - \int \tau_{h,j}^{sgs} n_j dS - \int \frac{\partial \bar{q}_i^r}{\partial x_i} dV \\ & + \int \left(\frac{\partial \bar{P}}{\partial t} + \tilde{u}_j \frac{\partial \bar{P}}{\partial x_j} \right) dV + \int \tau_{ij} \frac{\partial \tilde{u}_i}{\partial x_j} dV + \int S_\theta dV. \end{aligned} \quad (3.12)$$

The above equation is derived by starting with the total internal energy equation, subtracting the mechanical energy equation and enforcing the variable density continuity equation. Note that the above equation includes possible source terms due to thermal radiative transport, viscous dissipation, pressure work, and external driving sources (S_θ).

The simple Fickian diffusion velocity approximation, Equation (3.22), is assumed, so that the mean diffusive heat flux vector \bar{q}_j is

$$\bar{q}_j = - \left[\frac{\kappa}{C_p} \frac{\partial h}{\partial x_j} - \frac{\mu}{\text{Pr}} \sum_{k=1}^K h_k \frac{\partial Y_k}{\partial x_j} \right] - \frac{\mu}{\text{Sc}} \sum_{k=1}^K h_k \frac{\partial Y_k}{\partial x_j}.$$

If $\text{Sc} = \text{Pr}$, i.e., unity Lewis number ($\text{Le} = 1$), then the diffusive heat flux vector simplifies to $\bar{q}_j = -\frac{\mu}{\text{Pr}} \frac{\partial \tilde{h}}{\partial x_j}$. In the code base, the user has the ability to either specify a laminar Prandtl number, which is a constant, or provide a property evaluator for thermal conductivity. Inclusion of a Prandtl number prevails and ensures that the thermal conductivity is computed base on $\kappa = \frac{C_p \mu}{\text{Pr}}$. The viscous dissipation term is closed by

$$\begin{aligned} \tau_{ij} \frac{\partial u_i}{\partial x_j} &= \left((\mu + \mu_t) \left(\frac{\partial \tilde{u}_i}{\partial x_j} + \frac{\partial \tilde{u}_j}{\partial x_i} \right) - \frac{2}{3} \left(\tilde{\rho} \tilde{k} + \mu_t \frac{\partial \tilde{u}_k}{\partial x_k} \right) \delta_{ij} \right) \frac{\partial \tilde{u}_i}{\partial x_j} \\ &= \left[2\mu \tilde{S}_{ij} + 2\mu_t \left(\tilde{S}_{ij} - \frac{1}{3} \tilde{S}_{kk} \delta_{ij} \right) - \frac{2}{3} \tilde{\rho} \tilde{k} \delta_{ij} \right] \frac{\partial \tilde{u}_i}{\partial x_j}. \end{aligned}$$

The subgrid scale turbulent flux vector τ_h^{sgs} in Equation (3.12) is defined as

$$\tau_{hu_j} \equiv \bar{\rho} \left(\widetilde{hu_j} - \tilde{h} \tilde{u}_j \right).$$

As with species transport, the gradient diffusion hypothesis is used to close this subgrid scale model,

$$\tau_{h,j}^{sgs} = -\frac{\mu_t}{\text{Pr}_t} \frac{\partial \tilde{h}}{\partial x_j},$$

where Pr_t is the turbulent Prandtl number and μ_t is the modeled turbulent eddy viscosity from momentum closure. The resulting filtered and modeled turbulent energy equation is given by,

$$\begin{aligned} \int \frac{\partial \tilde{\rho} \tilde{h}}{\partial t} dV + \int \tilde{\rho} \tilde{h} \tilde{u}_j n_j dS &= \int \left(\frac{\mu}{\text{Pr}} + \frac{\mu_t}{\text{Pr}_t} \right) \frac{\partial \tilde{h}}{\partial x_j} n_j dS - \int \frac{\partial \tilde{q}_i^r}{\partial x_i} dV \\ &+ \int \left(\frac{\partial \tilde{P}}{\partial t} + \tilde{u}_j \frac{\partial \tilde{P}}{\partial x_j} \right) dV + \int \tau_{ij} \frac{\partial u_j}{\partial x_i} dV. \end{aligned} \quad (3.13)$$

The turbulent Prandtl number must have the same value as the turbulent Schmidt number for species transport to maintain unity Lewis number.

3.2.5 Review of Prandtl, Schmidt and Unity Lewis Number

For situations where a single diffusion coefficient is applicable (e.g., a binary gas system) the Lewis number is defined as:

$$\text{Le} = \frac{\text{Sc}}{\text{Pr}} = \frac{\alpha}{D}. \quad (3.14)$$

If the diffusion rates of energy and mass are equal,

$$\text{Sc} = \text{Pr} \text{ and } \text{Le} = 1. \quad (3.15)$$

For completeness, the thermal diffusivity, Prandtl and Schmidt number are defined by,

$$\alpha = \frac{\kappa}{\rho c_p}, \quad (3.16)$$

$$\text{Pr} = \frac{c_p \mu}{\kappa} = \frac{\mu}{\rho \alpha}, \quad (3.17)$$

and

$$\text{Sc} = \frac{\mu}{\rho D}, \quad (3.18)$$

where c_p is the specific heat, κ , is the thermal conductivity and α is the thermal diffusivity.

3.2.6 Thermal Heat Conduction

For non-isothermal object response that may occur in conjugate heat transfer applications, a simple single material heat conduction equation is supported.

$$\int \rho C_p \frac{\partial T}{\partial t} dV + \int q_j n_j dS = \int S dV. \quad (3.19)$$

where q_j is again the energy flux vector, however, now in the following temperature form:

$$q_j = -\kappa \frac{\partial T}{\partial x_j}.$$

3.2.7 ABL Forcing Source Terms

In LES of wind plant atmospheric flows, it is often necessary to drive the flow to a predetermined vertical velocity and/or temperature profile. In Nalu-Wind, this is achieved by adding appropriate source terms f_i to the momentum equation (3.1) and S_θ to the enthalpy equation (3.12).

First, the momentum source term is discussed. The main objective of this implementation is to force the volume averaged velocity at a certain location to a specified value ($\langle u_i \rangle = U_i$). The brackets used here, $\langle \rangle$, mean volume averaging over a certain region. In order to achieve this, a source term must be applied to the momentum equation. This source term can be better understood as a proportional controller within the momentum equation.

The velocity and density fields can be decomposed into a volume averaged component and fluctuations about that volume average as $u_i = \langle u_i \rangle + u'_i$ and $\bar{\rho} = \langle \bar{\rho} \rangle + \bar{\rho}'$. A decomposition of the plane averaged momentum at a given instance in time is then

$$\langle \bar{\rho} u_i \rangle = \langle \bar{\rho} \rangle \langle u_i \rangle + \langle \bar{\rho}' u'_i \rangle.$$

We now wish to apply a momentum source based on a desired spatial averaged velocity U_i . This can be expressed as:

$$\langle \bar{\rho} u_i^* \rangle = \langle \bar{\rho} \rangle \langle u_i^* \rangle + \langle \bar{\rho}' u_i^{*'} \rangle,$$

where u_i^* is an unknown reference velocity field whose volume average is the desired velocity $\langle u_i^* \rangle = U_i$. Since the correlation $\langle \bar{\rho}' u_i^{*'} \rangle$ is unknown, we assume that

$$\langle \bar{\rho}' u_i^{*'} \rangle = \langle \bar{\rho}' u'_i \rangle$$

such that the momentum source can now be defined as:

$$f_i = \alpha_u \left(\frac{\langle \bar{\rho} \rangle U_i - \langle \bar{\rho} \rangle \langle u_i \rangle}{\Delta t} \right) \quad (3.20)$$

where $\langle \rangle$ denotes volume averaging at a certain time t , U_i is the desired spatial averaged velocity, and Δt is the time-scale between when the source term is computed (time t) and when it is applied (time $t + \Delta t$). This is typically chosen to be the simulation time-step. In the case of an ABL simulation with flat terrain, the volume averaging is done over an infinitesimally thin slice in the x and y directions, such that the body force is only a function of height z and time t . The implementation allows the user to prescribe relaxation factors α_u for the source terms that are applied. Nalu-Wind uses a default value of 1.0 for the relaxation factors if no values are defined in the input file during initialization.

The enthalpy source term works similarly to the momentum source term. A temperature difference is computed at every time-step and a forcing term is added to the enthalpy equation:

$$S_\theta = \alpha_\theta C_p \left(\frac{\theta_{\text{ref}} - \langle \theta \rangle}{\Delta t} \right)$$

where θ_{ref} is the desired spatial averaged temperature, $\langle \theta \rangle$ is the spatial averaged temperature, C_p is the heat capacity, α_θ is the relaxation factor, and Δt is the time-scale.

The present implementation can vary the source terms as a function of time and space using either a user-defined table of previously computed source terms (e.g., from a *precursor* simulation or another model such as WRF), or compute the source term as a function of the transient flow solution.

3.2.8 Conservation of Species

The integral form of the Favre-filtered species equation used for turbulent transport is

$$\int \frac{\partial \tilde{\rho} \tilde{Y}_k}{\partial t} dV + \int \tilde{\rho} \tilde{Y}_k \tilde{u}_j n_j dS = - \int \tau_{Y_k, j}^{sgs} n_j dS - \int \overline{\rho Y_k \hat{u}_{j,k}} n_j dS + \int \bar{\omega}_k dV, \quad (3.21)$$

where the form of diffusion velocities (see Equation (3.22)) assumes the Fickian approximation with a constant value of diffusion velocity for consistency with the turbulent form of the energy equation, Equation (3.12). The simplest form is Fickian diffusion with the same value of mass diffusivity for all species,

$$\hat{u}_{j,k} = -D \frac{1}{Y_k} \frac{\partial Y_k}{\partial x_j}. \quad (3.22)$$

The subgrid scale turbulent diffusive flux vector $\tau_{Y_k, j}^{sgs}$ is defined as

$$\tau_{Y_k, j}^{sgs} \equiv \tilde{\rho} \left(\widetilde{Y_k u_j} - \tilde{Y}_k \tilde{u}_j \right).$$

The closure for this model takes on its usual gradient diffusion hypothesis, i.e.,

$$\tau_{Y_k, j}^{sgs} = - \frac{\mu_t}{Sc_t} \frac{\partial \tilde{Y}_k}{\partial x_j},$$

where Sc_t is the turbulent Schmidt number for all species and μ_t is the modeled turbulent eddy viscosity from momentum closure.

The Favre-filtered and modeled turbulent species transport equation is,

$$\int \frac{\partial \tilde{\rho} \tilde{Y}_k}{\partial t} dV + \int \tilde{\rho} \tilde{Y}_k \tilde{u}_j n_j dS = \int \left(\frac{\mu}{Sc} + \frac{\mu_t}{Sc_t} \right) \frac{\partial \tilde{Y}_k}{\partial x_j} n_j dS + \int \bar{\omega}_k dV. \quad (3.23)$$

If transporting both energy and species equations, the laminar Prandtl number must be equal to the laminar Schmidt number and the turbulent Prandtl number must be equal to the turbulent Schmidt number to maintain unity Lewis number. Although there is a species conservation equation for each species in a mixture of n species, only $n - 1$ species equations need to be solved since the mass fractions sum to unity and

$$\tilde{Y}_n = 1 - \sum_{j \neq n}^n \tilde{Y}_j.$$

Finally, the reaction rate source term is expected to be added based on an operator split approach whereby the set of ODEs are integrated over the full time step. The chemical kinetic source terms can be sub-integrated within a time step using a stiff ODE integrator package.

The following system of ODEs are numerically integrated over a time step Δt for a fixed temperature and pressure starting from the initial values of gas phase mass fraction and density,

$$\dot{Y}_k = \frac{\dot{\omega}_k(Y_k)}{\rho}.$$

The sources for the sub-integration are computed with the composition and density at the new time level which are used to approximate a mean production rate for the time step

$$\dot{\omega}_k \approx \frac{\rho^* Y_k^* - \rho Y_k}{\Delta t}.$$

3.2.9 Subgrid-Scale Kinetic Energy One-Equation LES Model

The subgrid scale kinetic energy one-equation turbulence model, or k^{sgs} model, [Dav97], represents a simple LES closure model. The transport equation for subgrid turbulent kinetic energy is given by

$$\int \frac{\partial \bar{\rho} k^{sgs}}{\partial t} dV + \int \bar{\rho} k^{sgs} \tilde{u}_j n_j dS = \int \frac{\mu_t}{\sigma_k} \frac{\partial k^{sgs}}{\partial x_j} n_j dS + \int (P_k^{sgs} - D_k^{sgs}) dV. \quad (3.24)$$

The production of subgrid turbulent kinetic energy, P_k^{sgs} , is modeled by,

$$P_k \equiv -\overline{\rho u_i'' u_j''} \frac{\partial \tilde{u}_i}{\partial x_j}, \quad (3.25)$$

while the dissipation of turbulent kinetic energy, D_k^{sgs} , is given by

$$D_k^{sgs} = \rho C_\epsilon \frac{k^{sgs \frac{3}{2}}}{\Delta},$$

where the grid filter length, Δ , is given in terms of the grid cell volume by

$$\Delta = V^{\frac{1}{3}}.$$

The subgrid turbulent eddy viscosity is then provided by

$$\mu_t = C_{\mu_\epsilon} \Delta k^{sgs \frac{1}{2}},$$

where the values of C_ϵ and C_{μ_ϵ} are 0.845 and 0.0856, respectively.

For simulations in which a buoyancy source term is desired, the code supports the Rodi form,

$$P_b = \beta \frac{\mu^T}{Pr} g_i \frac{\partial T}{\partial x_i}.$$

3.2.10 Shear Stress Transport (SST) RANS Model Suite

Although Nalu-Wind is primarily expected to be a LES simulation tool, RANS modeling is supported through the activation of the SST equation set.

It has been observed that standard 1998 $k - \omega$ models display a strong sensitivity to the free stream value of ω (see Menter, [MKL03]). To remedy, this, an alternative set of transport equations have been used that are based on smoothly blending the $k - \omega$ model near a wall with $k - \epsilon$ away from the wall. Because of the relationship between ω and ϵ , the transport equations for turbulent kinetic energy and dissipation can be transformed into equations involving k and ω . Aside from constants, the transport equation for k is unchanged. However, an additional cross-diffusion term is present in the ω equation. Blending is introduced by using smoothing which is a function of the distance from the wall, $F(y)$. The transport equations for the Menter 2003 model are then

$$\begin{aligned} \int \frac{\partial \bar{\rho} k}{\partial t} dV + \int \bar{\rho} k \tilde{u}_j n_j dS &= \int (\mu + \hat{\sigma}_k \mu_t) \frac{\partial k}{\partial x_j} n_j + \int (P_k^\omega - \beta^* \bar{\rho} k \omega) dV, \\ \int \frac{\partial \bar{\rho} \omega}{\partial t} dV + \int \bar{\rho} \omega \tilde{u}_j n_j dS &= \int (\mu + \hat{\sigma}_\omega \mu_t) \frac{\partial \omega}{\partial x_j} n_j + \int 2(1 - F) \frac{\bar{\rho} \sigma_\omega 2}{\omega} \frac{\partial k}{\partial x_j} \frac{\partial \omega}{\partial x_j} dV \\ &\quad + \int \left(\frac{\hat{\gamma}}{\nu_t} P_k^\omega - \hat{\beta} \bar{\rho} \omega^2 \right) dV. \end{aligned}$$

where the value of β^* is 0.09.

The model coefficients, $\hat{\sigma}_k$, $\hat{\sigma}_\omega$, $\hat{\gamma}$ and $\hat{\beta}$ must also be blended, which is represented by

$$\hat{\phi} = F \phi_1 + (1 - F) \phi_2.$$

where $\sigma_{k1} = 0.85$, $\sigma_{k2} = 1.0$, $\sigma_{\omega1} = 0.5$, $\sigma_{\omega2} = 0.856$, $\gamma_1 = \frac{5}{9}$, $\gamma_2 = 0.44$, $\beta_1 = 0.075$ and $\beta_2 = 0.0828$. The blending function is given by

$$F = \tanh(arg_1^4),$$

where

$$arg_1 = \min \left(\max \left(\frac{\sqrt{k}}{\beta^* \omega y}, \frac{500\mu}{\bar{\rho} y^2 \omega} \right), \frac{4\bar{\rho} \sigma_{\omega2} k}{CD_{k\omega} y^2} \right).$$

The final parameter is

$$CD_{k\omega} = \max \left(2\bar{\rho} \sigma_{\omega2} \frac{1}{\omega} \frac{\partial k}{\partial x_j} \frac{\partial \omega}{\partial x_j}, 10^{-10} \right).$$

An important component of the SST model is the different expression used for the turbulent viscosity,

$$\mu_t = \frac{a_1 \bar{\rho} k}{\max(a_1 \omega, SF_2)},$$

where F_2 is another blending function given by

$$F_2 = \tanh(arg_2^2).$$

The final parameter is

$$arg_2 = \max \left(\frac{2\sqrt{k}}{\beta^* \omega y}, \frac{500\mu}{\bar{\rho} \omega y^2} \right).$$

3.2.11 Detached Eddy Simulation (DES) Formulation

The DES technique is also supported in the code base when the SST model is activated. This model seeks to formally relax the RANS-based approach and allows for a theoretical basis to allow for transient flows. The method follows the method of Temporally Filtered NS formulation as described by Tieszen, [TDB05].

The SST DES model simply changes the turbulent kinetic energy equation to include a new minimum scale that manipulates the dissipation term.

$$D_k = \frac{\rho k^{3/2}}{l_{DES}},$$

where l_{DES} is the $\min(l_{SST}, c_{DES} l_{DES})$. The constants are given by, $l_{SST} = \frac{k^{1/2}}{\beta^* \omega}$ and c_{DES} represents a blended set of DES constants: $c_{DES_1} = 0.78$ and $c_{DES_2} = 0.61$. The length scale, l_{DES} is the maximum edge length scale touching a given node.

3.2.12 Active Model Split (AMS) Formulation

The AMS approach is a recently developed hybrid RANS/LES framework by Haering, Oliver and Moser [HOM20]. In this approach a triple decomposition is used, breaking the instantaneous velocity field into an average component $\langle u_i \rangle$, a resolved fluctuation $u_i^>$ and an unresolved fluctuation $u_i^<$. The subgrid stress is split into two terms, $\tau_{ij} = \tau_{ij}^{SGRS} + \tau_{ij}^{SGET}$, with τ_{ij}^{SGRS} modeling the mean subgrid stress, taking on the form of a standard RANS subgrid stress model and τ_{ij}^{SGET} representing the energy transfer from the resolved to the modeled scales. In addition,

a forcing stress F_i is added to the momentum equations to induce the transfer of energy from the modeled to the resolved scales. Thus the AMS approach solves the following momentum equation

$$\begin{aligned} & \int \frac{\partial \bar{\rho} \tilde{u}_i}{\partial t} dV + \int \bar{\rho} \tilde{u}_i \tilde{u}_j n_j dS + \int \left(\bar{P} + \frac{2}{3} \bar{\rho} k \right) n_i dS = \\ & \int 2\mu \left(\tilde{S}_{ij} - \frac{1}{3} \tilde{S}_{kk} \delta_{ij} \right) n_j dS + \int (\bar{\rho} - \rho_o) g_i dV + \\ & \int 2\mu_t \left(\langle S_{ij} \rangle - \frac{1}{3} \langle S_{kk} \rangle \delta_{ij} \right) n_j dS + \int (\mu_{ik}^t \tilde{u}_j + \mu_{jk}^t \tilde{u}_i) n_k dS + \int f_i dV. \end{aligned} \quad (3.26)$$

Split subgrid model stress

In a typical Hybrid RANS/LES approach, the observation that the LES and RANS equations take on the same mathematical form is leveraged, relying simply on a modified turbulent viscosity coefficient that takes into account the ability to resolve some turbulent content. Due to deficiencies in this approach as discussed in Haering et al. [HOM20], an alternative approach where the modeled term is split into two contributions, one representing the impact on the mean flow and one the impact on the resolved fluctuations, from the unresolved content, is used in the Active Model Split (AMS) formulation.

Starting by substitution of a triple decomposition of the flow variables, $\phi = \langle \phi \rangle + \phi^> + \phi^<$, with $\langle \cdot \rangle$ representing a mean quantity, $\phi^>$ a resolved fluctuation and $\phi^<$ an unresolved fluctuation and dropping all terms that have an unresolved fluctuation in them (since by definition, these terms cannot be resolved and thus must be modeled) we get the following momentum equation:

$$\frac{\partial \bar{u}_i}{\partial t} + \frac{\partial \bar{u}_i \bar{u}_j}{\partial x_j} = -\frac{1}{\rho} \left(\frac{\partial \bar{P}}{\partial x_i} \right) + \nu \frac{\partial^2 \bar{u}_i}{\partial x_j \partial x_j} + \frac{\partial \tau_{ij}^M}{\partial x_j} + F_i$$

Note that here, $\bar{\phi} = \langle \phi \rangle + \phi^>$ represents an instantaneous resolved quantity and F_i is a forcing term discussed in Sec. [AMS forcing](#).

The model term in AMS, τ_{ij}^M is split into two pieces, the first representing the impact of the unresolved scales on the mean flow, referred to as τ_{ij}^{SGRS} , since this mimics the purpose of RANS models and seeks to model the subgrid Reynolds Stress (SGRS). The second term represents the impact of the unresolved scales on the resolved fluctuations which acts to capture energy transfer from the resolved fluctuations to the unresolved fluctuations, which as Haering et al. points out, is really the primary function of typical LES SGS models. As this term models subgrid energy transfer (SGET), it is referred to as τ_{ij}^{SGET} .

τ_{ij}^{SGRS} is modeled using a typical RANS model, but since in the hybrid context, some turbulence is resolved, the magnitude of the stress tensor is reduced through a derived scaling with $\alpha = \beta^{1.7}$, $\beta = 1 - k_{res}/k_{tot}$, where k_{tot} is the total kinetic energy, taken from the RANS model and $k_{res} = 0.5 \langle u_i^> u_i^> \rangle$, a measure of the average resolved turbulent kinetic energy.

τ_{ij}^{SGET} is modeled using the M43 SGS model discussed in Haering et al. [HOM19]. This uses an anisotropic representation, $\tau_{ij} = \nu_{jk} \partial_k u_i + \nu_{ik} \partial_k u_j$, of the stress tensor and a tensorial eddy viscosity, $\nu_{ij} = C(\mathcal{M}) \langle \epsilon \rangle^{1/3} \mathcal{M}_{ij}^{4/3}$, with C , a coefficient determined as a function of the eigenvalues of \mathcal{M} , a metric tensor measure of the grid and $\langle \epsilon \rangle$ inferred from the RANS model.

So this produces the final form for the AMS model term,

$$\begin{aligned} \tau_{ij}^M &= \tau_{ij}^{SGRS} + \tau_{ij}^{SGET} \\ &= 2\alpha(2 - \alpha) \nu_t^{RANS} \langle S_{ij} \rangle - \frac{2}{3} \beta k_{tot} \delta_{ij} + C(\mathcal{M}) \langle \epsilon \rangle^{1/3} \left(\mathcal{M}_{jk}^{4/3} \frac{\partial u_i^>}{\partial x_k} + \mathcal{M}_{ik}^{4/3} \frac{\partial u_j^>}{\partial x_k} \right). \end{aligned}$$

The AMS model terms are implemented for the edge based scheme in *MomentumSSTAMSDiffEdgeKernel*. The isotropic component, $2\beta k_{tot} \delta_{ij}/3$ is brought into the pressure term.

Averaging functions

The averaging algorithms are invoked as part of the *AMSAvgDriver* and are called from the *pre_iter_work* function so that they are executed at the beginning of each Picard iteration. The *AMSAvgDriver* is invoked last, so to ensure that this is also done initially, so that the initial step has correct average quantities, the averaging functions are also called in the *initial_work* function.

The main averaging algorithm is *SSTAMSAveragesAlg*. The averaging function is solving a simple causal average equation for the intermediate (or final) quantity:

$$\frac{\partial \langle \phi^* \rangle}{\partial t} = \frac{1}{T_{RANS}^*} (\phi^* - \langle \phi^n \rangle)$$

Here $\langle \cdot \rangle$ refers to a mean (time-averaged) quantity and T_{RANS} is the timescale of the turbulence determined by the underlying RANS scalars ($1/(\beta^* \omega)$ in SST). $(\cdot)^n$ refers to a previous timestep quantity and $(\cdot)^*$ refers to an intermediate quantity. Note that currently the time scale is stored in a nodal field.

We can discretize the causal average equation explicitly to produce the implemented form:

$$\begin{aligned} \langle \phi^* \rangle &= \langle \phi^n \rangle + \frac{\Delta t}{T_{RANS}^*} (\phi^* - \langle \phi^n \rangle) \\ \langle \phi^* \rangle &= \frac{\Delta t}{T_{RANS}^*} \phi^* + \left(1 - \frac{\Delta t}{T_{RANS}^*} \right) \langle \phi^n \rangle \end{aligned}$$

The averaging is carried out for velocities (u_i), velocity gradients ($\partial u_i / \partial x_j$), pressure (P), density (ρ), resolved turbulent kinetic energy ($k_{res} = 0.5 u_i^> u_i^>$) and the kinetic energy production ($\mathcal{P}_k = \langle S_{ij} \rangle (\tau_{ij}^{SGRS} - u_i^> u_j^>)$). Note that $>$ is used to denote a resolved fluctuation, i.e. $\phi^> = \phi - \langle \phi \rangle$.

Dynamic Hybrid Diagnostics

Typically in a hybrid model, it is necessary to have diagnostics that assess the ability of the grid to resolve turbulent content and to aid in its introduction. In AMS, there are two main diagnostic quantities, $\alpha = \beta^{1.7} = (1 - k_{res}/k_{tot})^{1.7}$ and a resolution adequacy parameter, r_k , which is used to evaluate where in the flow the grid and the amount of resolved turbulent content is inconsistent.

β is a straight-forward calculation. Limiters are applied to β to realize the RANS and DNS limits. In the RANS limit, $k_{res} = 0$ and thus $\beta = 1$, so β is limited from evaluation above 1. In the DNS limit, ideally, the ratio of the approximate Kolmogorov velocity scale to total TKE would be used as a lower bound,

$$\beta = \max \left(1 - \frac{k_{res}}{k_{tot}}, \frac{(\nu \epsilon)^{1/2}}{k_{tot}} \right),$$

but that has shown some performance issues near the wall when using SST with AMS. Currently an adhoc lower bound of $\beta = 0.01$ is used instead. The resolution adequacy parameter is based on the ratio between the anisotropic grid metric tensor, $\mathcal{M} = \mathcal{J}^T J$, where \mathcal{J} is the mapping from a unit cube to the cell, and the length scale associated with the model production. It takes the form,

$$r_k = \left(\frac{3}{2 \langle \bar{v}^2 \rangle} \right)^{3/2} \max_{\lambda} (\mathcal{P}_{ik}^{SGS} \mathcal{M}_{kj}).$$

For the RANS models used in Nalu-Wind, $\langle \bar{v}^2 \rangle \approx 5 \nu_{RANS} / T_{RANS}$. $\mathcal{P}_{ij}^{SGS} = \frac{1}{2} (\tau_{ik} \partial \bar{u}_k / \partial x_j + \tau_{jk} \partial \bar{u}_k / \partial x_i)$ is the full subgrid production tensor, with $\tau_{ij} = \tau_{ij}^{SGRS} + \tau_{ij}^{SGET} + \frac{2}{3} \alpha k_{tot} \delta_{ij}$.

Forcing Term

When the grid is capable of resolving some turbulent content, the model will want to reduce the modeled stress and allow for resolved turbulence to contribute the remaining piece of the total stress. As discussed in Haering et al. [HOM20] and the observation of “modeled-stress depletion” in other hybrid approaches, such as DES, a mechanism for inducing resolved turbulent fluctuations at proper energy levels and timescales to match your reduction of the modeled stress is needed. AMS resolves this issue through the use of an active forcing term, designed to introduce turbulent fluctuations into regions of the grid where turbulent content can be supported. The implications of the specific form and method of introduction for this forcing term is still an area of ongoing research, but for now, empirical testing has shown great success with a simple turbulent structure based off of Taylor-Green vortices.

The forcing term F_i is determined by first specifying an auxiliary field based off of a Taylor-Green vortex:

$$\begin{aligned} h_1 &= \frac{1}{3} \cos(a_x x'_1) \sin(a_y x'_2) \sin(a_z x'_3), \\ h_2 &= -\sin(a_x x'_1) \cos(a_y x'_2) \sin(a_z x'_3), \\ h_3 &= \frac{2}{3} \sin(a_x x'_1) \sin(a_y x'_2) \cos(a_z x'_3), \end{aligned}$$

with $\mathbf{x}' = \mathbf{x} + \langle \mathbf{u} \rangle t$ and $a_i = \pi / \mathbb{P}_i$. \mathbb{P} is determined as follows,

$$\begin{aligned} l &= \frac{4 - (1 - \max(\beta, 0.8)) (\beta k)^{3/2}}{0.4 \epsilon} \\ l &= \min \left(\max \left(l, 70 \frac{\nu^{3/4}}{\epsilon^{1/4}} \right), d \right) \\ l'_i &= \min(l, L_{p_i}) \\ f_i &= \text{nint} \left(\frac{L_{p_i}}{l'_i} \right) \\ \mathbb{P}_i &= \frac{L_{p_i}}{f_i}, \end{aligned}$$

where L_{p_i} is related to the periodic domain size and is user specified. With the initial TG vortex field, h_i , determined,

we now determine a scaling factor (η) for the forcing.

$$\begin{aligned}
 T_\beta &= \max\left(\beta k / \epsilon, 6\sqrt{\nu/\epsilon}\right) \\
 F_{tar} &= 8\sqrt{\alpha \bar{v}^2} / T_\beta \\
 \mathcal{P}_r &= \Delta t F_{tar} (h_i u_i^{\geq}) \\
 \beta_K &= \min(\sqrt{\nu \epsilon / k}, 1) \\
 \hat{\beta} &= \begin{cases} \frac{1 - \beta}{1 - \beta_K} & \beta_K < 1 \\ 10000 & \text{else} \end{cases} \\
 C_f &= -1 \tanh\left(1 - \frac{1}{\sqrt{\min(\langle r_k \rangle, 1)}}\right) \\
 C_f &= C_f(1.0 - \min(\tanh(10(\hat{\beta} - 1)) + 1, 1)) \\
 \eta &= \begin{cases} F_{tar} C_f & \langle r_k \rangle < 1, \mathcal{P}_r \geq 0 \\ 0 & \text{else} \end{cases}
 \end{aligned}$$

Now the final forcing field, $F_i = \eta h_i$. Since this is being added as a source term to the momentum solve, we are not projecting onto a divergence free field and are instead allowing that to pass into the continuity solve, where the intermediate velocity field with the forcing will then be projected onto a divergence free field. This is implemented in the node kernels as *MomentumSSTAMSForsingNodeKernel*.

3.2.13 Solid Stress

A fully implicit CVFEM (only) linear elastic equation is supported in the code base. This equation is either used for true solid stress prediction or for smoothing the mesh due to boundary mesh motion (either through fluid structure interaction (FSI) or prescribed mesh motion).

Consider the displacement for component i , u_i equation set,

$$\rho \frac{\partial^2 u_i}{\partial t^2} - \frac{\partial \sigma_{ij}}{\partial x_j} = F_i, \quad (3.27)$$

where the Cauchy stress tensor, σ_{ij} assuming Hooke's law is given by,

$$\sigma_{ij} = \mu \left(\frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i} \right) + \lambda \frac{\partial u_k}{\partial x_k} \delta_{ij}. \quad (3.28)$$

Above, the so-called Lamé coefficients, Lamé's first parameter, λ (also known as the Lamé modulus) and Lamé's second parameter, μ (also known as the shear modulus) are provided as functions of the Young's modulus, E , and Poisson's ratio, ν ; here shown in the context of an isotropic elastic material,

$$\mu = \frac{E}{2(1 + \nu)}, \quad (3.29)$$

and

$$\lambda = \frac{E\nu}{(1 + \nu)(1 - 2\nu)}. \quad (3.30)$$

Note that the above notation of u_i to represent displacement is with respect to the classic definition of current and model coordinates,

$$x_i = X_i + u_i \quad (3.31)$$

where x_i is the position, relative to the fixed, or previous position, X_i .

The above equations are solved for mesh displacements, u_i . The supplemental relationship for solid velocity, v_i is given by,

$$v_i = \frac{\partial u_i}{\partial t}. \quad (3.32)$$

Numerically, the velocity might be obtained by a backward Euler or BDF2 scheme,

$$v_i = \frac{\gamma_1 u_i^{n+1} + \gamma_2 u_i^n + \gamma_3 u_i^{n-1}}{\Delta t} \quad (3.33)$$

3.2.14 Moving Mesh

The code base supports three notions of moving mesh: 1) linear elastic equation system that computes the stress of a solid 2) solid body rotation mesh motion and 3) mesh deformation via an external source.

The linear elastic equation system is activated via the standard equation system approach. Properties for the solid are specified in the material block. Mesh motion is prescribed by the input file via the `mesh_motion` block. Here, it is assumed that the mesh motion is solid rotation. For fluid/structure interaction (FSI) a mesh smoothing scheme is used to propagate the surface mesh displacement obtained by the solids solve. Simple mesh smoothing is obtained via a linear elastic solve in which the so-called Lamé constants are proportional to the inverse of the dual volume. This allows for boundary layer mesh locations to be stiff while free stream mesh elements to be soft.

Additional mesh motion terms are required for the Eulerian fluid mechanics solve. Using the geometric conservative law the time and advection source term for a general scalar ϕ can be written as:

$$\int \frac{\rho \phi}{\partial t} dV + \int \rho \phi (u_j - v_j) n_j dS + \int \rho \phi \frac{\partial v_k}{\partial x_j} dV, \quad (3.34)$$

where u_j is the fluid velocity and v_j is the mesh velocity. Mesh velocities and the mesh velocity spatial derivatives are provided by the mesh smoothing solve. Activating the external mesh deformation or mesh motion block will result in the velocity relative to mesh calculation in the advection terms. The line command for source term, “*gcl*” must be activated for each equation for the volume integral to be included in the set of PDE solves. Finally, transfers are expected between the physics. For example, the solids solve is to provide mesh displacements to the mesh smoothing realm. The mesh smoothing procedure provides the boundary velocity, mesh velocity and projected nodal gradients of the mesh velocity to the fluids realm. Finally, the fluids solve is to provide the surface force at the desired solids surface. Currently, the pressure is transferred from the fluids realm to the solids realm. The ideal view of FSI is to solve the solids pde at the half time step. As such, in time, the $P^{n+\frac{1}{2}}$ is expected. The `fsi_interface` input line command attribute is expected to be set at these unique surfaces. More advanced FSI coupling techniques are under development by a current academic partner.

3.2.15 Radiative Transport Equation

The spatial variation of the radiative intensity corresponding to a given direction and at a given wavelength within a radiatively participating material, $I(s)$, is governed by the Boltzmann transport equation. In general, the Boltzmann equation represents a balance between absorption, emission, out-scattering, and in-scattering of radiation at a point. For combustion applications, however, the steady form of the Boltzmann equation is appropriate since the transient term only becomes important on nanosecond time scales which is orders of magnitude shorter than the fastest chemical.

Experimental data shows that the radiative properties for heavily sooting, fuel-rich hydrocarbon diffusion flames ($10^{-4}\%$ to $10^{-6}\%$ soot by volume) are dominated by the soot phase and to a lesser extent by the gas phase. Since soot emits and absorbs radiation in a relatively constant spectrum, it is common to ignore wavelength effects when modeling radiative transport in these environments. Additionally, scattering from soot particles commonly generated by hydrocarbon flames is several orders of magnitude smaller than the absorption effect and may be neglected. Moreover, the phase function is rarely known. However, for situations in which the phase function can be approximated by the iso-tropic scattering assumption, i.e., an intensity for direction k has equal probability to be scattered in any direction l , the appropriate form of the Boltzmann radiative transport is

$$s_i \frac{\partial}{\partial x_i} I(s) + (\mu_a + \mu_s) I(s) = \frac{\mu_a \sigma T^4}{\pi} + \frac{\mu_s}{4\pi} G, \quad (3.35)$$

where μ_a is the absorption coefficient, μ_s is the scattering coefficient, $I(s)$ is the intensity along the direction s_i , T is the temperature and the scalar flux is G . The black body radiation, I_b , is defined by $\frac{\sigma T^4}{\pi}$. Note that for situations in which the scattering coefficient is zero, the RTE reduces to a set of linear, decoupled equations for each intensity to be solved.

The flux divergence may be written as a difference between the radiative emission and mean incident radiation at a point,

$$\frac{\partial q_i^r}{\partial x_i} = \mu_a [4\sigma T^4 - G], \quad (3.36)$$

where G is again the scalar flux. The flux divergence term is the same regardless of whether or not scattering is active. The quantity, $G/4\pi$, is often referred to as the mean incident intensity. Note that when the scattering coefficient is non-zero, the RTE is coupled over all intensity directions by the scalar flux relationship.

The scalar flux and radiative flux vector represent angular moments of the directional radiative intensity at a point,

$$G = \int_0^{2\pi} \int_0^\pi I(s) \sin \theta_{zn} d\theta_{zn} d\theta_{az},$$

$$q_i^r = \int_0^{2\pi} \int_0^\pi I(s) s_i \sin \theta_{zn} d\theta_{zn} d\theta_{az},$$

where θ_{zn} and θ_{az} are the zenith and azimuthal angles respectively as shown in Figure Fig. 3.1.

The radiation intensity must be defined at all portions of the boundary along which $s_i n_i < 0$, where n_i is the outward directed unit normal vector at the surface. The intensity is applied as a weak flux boundary condition which is determined from the surface properties and temperature. The diffuse surface assumption provides reasonable accuracy for many engineering combustion applications. The intensity leaving a diffuse surface in all directions is given by

$$I(s) = \frac{1}{\pi} [\tau \sigma T_\infty^4 + \epsilon \sigma T_w^4 + (1 - \epsilon - \tau) K], \quad (3.37)$$

where ϵ is the total normal emissivity of the surface, τ is the transmissivity of the surface, T_w is the temperature of the boundary, T_∞ is the environmental temperature and H is the incident radiation, or irradiation (incoming radiative flux). Recall that the relationship given by Kirchoff's Law that relates emissivity, transmissivity and reflectivity, ρ , is

$$\rho + \tau + \epsilon = 1.$$

where it is implied that $\alpha = \epsilon$.

3.2.16 Wall Distance Computation

RANS and DES simulations using $k - \omega$ *SST* or *SST-DES* equations requires the specification of a *wall distance* for computing various turbulence parameters. For static mesh simulations this field can be generated using a pre-processing step and provided as an input in the mesh database. However, for moving mesh simulations, e.g., blade




Fig. 3.1: Ordinate Direction Definition, $\mathbf{s} = \sin \theta_{zn} \sin \theta_{az} \mathbf{i} + \cos \theta_{zn} \mathbf{j} + \sin \theta_{zn} \cos \theta_{az} \mathbf{k}$.

resolved wind turbine simulations, this field must be computed throughout the course of the simulation. Nalu-Wind implements a Poisson equation ([G03]) to determine the wall distance d using the gradients of a field ϕ .

$$\nabla^2 \phi = 1$$

$$d = \pm \sqrt{\sum_{j=1,3} \left(\frac{\partial \phi}{\partial x_j} \right)^2} + \sqrt{\sum_{j=1,3} \left(\frac{\partial \phi}{\partial x_j} \right)^2 + 2\phi}$$

3.3 Discretization Approach

Nalu-Wind supports two discretizations: control volume finite element and (CVFEM) edge-based vertex centered (EBVC). Each are finite volume formulations and each solve for the primitives are each considered vertex-based schemes. Considerable testing has provided a set of general rules as to which scheme is optimal. In general, all equations and boundary conditions support either equation discretization with exception of the solid stress equation which has only been implemented for the CVFEM technique.

For generalized unstructured meshes that have poor quality, CVFEM has been shown to excel in accuracy and robustness. This is mostly due to the inherent accuracy limitation for the non-orthogonal correction terms that appear in the diffusion term and pressure stabilization for the EBVC scheme. For generalized unstructured meshes of decent quality, either scheme is ideal. Finally, for highly structured meshes with substantial aspect ratios, the edge-based scheme is ideal.

In general, the edge-based scheme is at least two times faster per iteration than the element-based scheme. For some classes of flows, it can be up to four times faster. However, due to the lagged coupling between the projected nodal gradient equation and the dofs, on meshes with high non-orthogonality, nonlinear residual convergence can be delayed.

3.3.1 CVFEM Dual Mesh

The classic low Mach algorithm uses the finite volume technique known as the control volume finite element method, see Schneider, [SR87], or Domino, [Dom06]. Control volumes (the mesh dual) are constructed about the nodes, shown in Figure Fig. 3.2 (upper left). Each element contains a set of sub-faces that define control-volume surfaces. The sub-faces consist of line segments (2D) or surfaces (3D). The 2D segments are connected between the element centroid and the edge centroids. The 3D surfaces (not shown here) are connected between the element centroid, the element face centroids, and the edge centroids. Integration points also exist within the sub-control volume centroids.

Recent work by Domino, [Dom14], has provided a proof-of-concept higher order CVFEM implementation whereby the linear basis and dual mesh definition is extended to higher order. The current code base supports the usage of P=2 elements (quadratic) for both 2D and 3D quad/hex topologies. This method has been formally demonstrated to be third-order spatially accurate and second-order in-time accurate. General polynomial promotion has been deployed in the higher order github branch. Figure Fig. 3.2 illustrates a general polynomial promotion from P=1 to P=6 and demonstrated spectral convergence using the method of manufactured solutions in Figure Fig. 3.3.

When using CVFEM, the discretized equations described in this manual are evaluated at either subcontrol-surface integration points (terms that have been integrated by parts) or at the subcontrol volume (time and source terms). Interpolation within the element is obtained by the standard elemental basis functions,

$$\phi_{ip} = \sum N_k^{ip} \phi_k. \quad (3.38)$$

where the index k represents a loop over all nodes in the element.

Gradients at the subcontrol volume surfaces are obtained by taking the derivative of Eq. (3.38), to obtain,

$$\frac{\partial \phi_{ip}}{\partial x_j} = \sum \frac{\partial N_{j,k}^{ip}}{\partial x_j} \phi_k. \quad (3.39)$$

The usage of the CVFEM methods results in the canonical 27-point stencil for a structured hexahedral mesh.

3.3.2 Edge-Based Discretization

In the edge-based discretization, the dual mesh defined in the CVFEM method is used to pre-process both dual mesh nodal volumes (needed in source and time terms) and edge-based area vectors (required for integrated-by-parts quantities, e.g., advection and diffusion terms).

Consider Figure Fig. 3.4, which is the original set of CVFEM dual mesh quadrature points shown above in Figure Fig. 3.2. Specifically, there are four subcontrol volumes about node 5 that contribute to the nodal volume dual mesh. In an edge-based scheme, the time and source terms use single point quadrature by assembling these four subcontrol volume contributions (eight in 3D) into one single nodal volume. In most cases, source terms may include gradients that are obtained by using the larger element-based stencil.

The same reduction of gauss points is realized for the area vector. Consider the edge between nodes 5 and 6. In the full CVFEM approach, subcontrol surfaces within the top element (5,6,9,8) and bottom element (2,3,6,5) are reduced to a single area vector at the edge midpoint of nodes 5 and 6. Therefore, advection and diffusion is now done in a manner very consistent with a cell centered scheme, i.e., classic “left”/“right” states.

The consolidation of time and source terms to nodal locations along with advection and diffusion at the edge midpoint results in a canonical five-point stencil in 2D and seven in 3D. Note the ability to handle hybrid meshes is readily performed one nodal volume and edge area are pre-processed. Edges and nodes are the sole topology that are iterated, thus making this scheme highly efficient, although inherently limited to second order spatial order of accuracy.

In general, the edge-based scheme is second order spatially accurate. Formal verification has been done to evaluate the accuracy of the EBVC relative to other implemented methods (Domino, [Dom14]). The edge-based scheme, which is based on dual mesh post-processing, represents a commonly used finite volume method in gas dynamics applications. The method also lends itself to psuedo-higher order methodologies by the blending of extrapolated values using the




Fig. 3.2: Polynomial promotion for a canonical CVPFEM quad element patch from $P = 1$ to $P = 6$.




Fig. 3.3: A recent spectral convergence plot using the Method of Manufactured Solutions for $P = 1$ through $P = 8$.




Fig. 3.4: A control volume centered about a finite-element node in a collection of 2-D quadrilateral elements (from [Dom06].)

projected nodal gradient and gauss point values (as does CVFEM). This provides a fourth order accurate diffusion and advection operator on a structured mesh.

The use of a consistent mass matrix is less apparent in edge-based schemes. However, if desired, the full element-based stencil can be used by iterating elements and assembling to the nodes.

The advantage of edge-based schemes over cell centered schemes is that the scheme naturally allows for a mixed elemental discretization. Projected nodal gradients can be element- or edge-based. LES filters and nodal gradients can also exploit the inherent elemental basis that exists in the pure CVFEM approach. In our experience, the optimal scheme on high quality meshes uses the CVFEM for the continuity solve and EBVC discretization for all other equations. This combination allows for the full CVFEM diffusion operator for the pressure Poisson equation and the EBVC approach for equations where inverse Reynolds scaling reduces the importance of the diffusion operator. This scheme can be activated by the use of the `use_edges: yes` Realm line command in combination of the `LowMachEOM` system line command, `element_continuity_eqs: yes`.

3.3.3 Projected Nodal Gradients

In the edge or element-based algorithm, projected nodal gradients are commonplace. Projected nodal gradients are used in the fourth order pressure stabilization terms, higher order upwind methods, discontinuity capturing operators (DCO) and turbulence source terms. For an edge-based scheme, they are also used in the diffusion term when non-orthogonality of the mesh is noted.

There are many procedures for determining the projected nodal gradient ranging from element-based schemes to edge-based approaches. In general, the projected nodal gradient is viewed as an L_2 minimization between the discontinuous Gauss-point value and the continuous nodal value. The projected nodal gradient, in an L_2 sense is given by,

$$\int w G_j \phi dV = \int \frac{\partial \phi}{\partial x_j} dV. \quad (3.40)$$

Using integration-by-parts and a piece-wise constant test function, the above equation is written as,

$$\int w_I G_j \phi dV = \int \phi_{ip} n_j dS. \quad (3.41)$$

For a lumped L_2 projected nodal gradient, the approach is based on a Green-Gauss integration,

$$G_j \phi = \frac{\int \phi_{ip} A_j}{dV}. \quad (3.42)$$

In the above lumped mass matrix approach, the value at the integration point can either be based on the CVFEM dual mesh evaluated at the subcontrol surface, i.e., the line command option, `element` or the `edge`, which evaluates the term at the edge midpoint using the assemble edge area vector. In all cases, the lumped mass matrix approach is strictly second order accurate. When running higher order CVFEM, a consistent mass matrix approach is required to maintain design order of the overall discretization. This is strictly due to the pressure stabilization whose accuracy can be affected by the form of the projected nodal gradient (see the Nalu-Wind theory manual or a variety of SNL-based publications).

In the description that follows, $\bar{G}_j \phi$ represent the average nodal gradient evaluated at the integration point of interest.

The choice of projected nodal gradients is specified in the input file for each dof. Keywords `element` or `edge` are used to define the form of the projection. The forms of the projected nodal gradients is arbitrary relative to the choosed underlying discretization. For strongly non-orthogonal meshes, it is recommended to use an element-based projected nodal gradient for the continuity equation when the EBVC method is in use. In some limited cases, e.g., pressure, mixture fraction and enthalpy, the `manage-png` line command can be used to solve the simple linear system for the consistent mass matrix.

3.3.4 Time and Source Terms

Time and source terms also volumetric contributions and also use the dual nodal volume. In both discretization approaches, this assembly is achieved as a simple nodal loop. In some cases, e.g., the k_{sgs} partial differential equation, the source term can use projected nodal gradients.

$$\int \frac{\partial \rho \phi}{\partial t} dV = \int S_\phi dV$$

3.3.5 Diffusion

As already noted, for the CVFEM method, the diffusion term at the subcontrol surface integration points use the the elemental shape functions and derivatives. For the standard diffusion term, and using Eq. (3.39), the CVFEM diffusion operator contribution at a given integration point (here simply demonstrated for a 2D edge with prescribed area vector) is as follows,

$$-\int \Gamma \frac{\partial \phi}{\partial x_j} A_j = -\Gamma_{ip} \left[\left(\frac{\partial N_0^{ip}}{\partial x} \phi_0 + \frac{\partial N_1^{ip}}{\partial x} \phi_1 \right) A_x + \left(\frac{\partial N_0^{ip}}{\partial y} \phi_0 + \frac{\partial N_1^{ip}}{\partial y} \phi_1 \right) A_y \right]$$

Standard Gauss point locations at the subcontrol surfaces can be shifted to the edge-midpoints for a more stable (monotonic) diffusion operator that is better conditioned for high aspect ratio meshes.

For the edge-based diffusion operator, special care is noted as there is no ability to use the elemental basis to define the diffusion operator. As with cell-centered schemes, non-orthogonal contributions for the diffusion operator arise due to a difference in direction between the assembled edge area vector and the distance vector between nodes on an edge. On skewed meshes, this non-orthogonality can not be ignored.

Following the work of Jasek, [Jas96], the over-relaxed approach is used. The form of any gradient for direction j for field ϕ is

$$\frac{\partial \phi}{\partial x_j}_{ip} = G_j^- \phi + [(\phi_R - \phi_L) - G_L^- \phi dx_L] \frac{A_j}{A_k dx_k}. \quad (3.43)$$

In the above expression, we are iterating edges with a Left node L and Right node R along with edge-area vector, A_j . The $G_j^- \phi$ is simple averaging of the left and right nodes to the edge midpoint. In general, a standard edge-based diffusion term is written as,

$$-\int \Gamma \frac{\partial \phi}{\partial x_j} A_j = -\Gamma_{ip} \left[(G_x^- \phi A_x + G_y^- \phi A_y) + (\phi_R - \phi_L) \frac{A_x A_x + A_y A_y}{A_x dx_x + A_y dx_y} - (G_x^- \phi dx_x + G_y^- \phi dx_y) \frac{A_x A_x + A_y A_y}{A_x dx_x + A_y dx_y} \right].$$

Momentum Stress

The viscous stress tensor, τ_{ij} is formed based on the standard gradients defined above for either the edge or element-based discretization. The viscous force for component i is given by,

$$-\int \tau_{ij} A_j = -\int \mu_{ip} \left(\frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i} \right) A_j.$$

For example, the x and y-component of viscous force is given by,

$$F_x = -\mu_{ip} \left(\frac{\partial u_x}{\partial x} A_x + \frac{\partial u_x}{\partial y} A_y \right) - \mu_{ip} \left(\frac{\partial u_x}{\partial x} A_x + \frac{\partial u_y}{\partial x} A_y \right),$$

$$F_y = -\mu_{ip} \left(\frac{\partial u_y}{\partial x} A_x + \frac{\partial u_y}{\partial y} A_y \right) - \mu_{ip} \left(\frac{\partial u_x}{\partial y} A_x + \frac{\partial u_y}{\partial y} A_y \right).$$

Note that the first part of the viscous stress is simply the standard diffusion term. Note that the so-called non-solenoidal viscous stress contribution is frequently written in terms of projected nodal gradients. However, for CVFEM this procedure is rarely used given the elemental basis definition. As such, the use of shape function derivatives is clear.

The viscous stress contribution at an integration point for CVFEM (again using the 2D example with variable area vector) can be written as,

$$\begin{aligned} F_x &= -\Gamma_{ip} \left[\left(\frac{\partial N_0^{ip}}{\partial x} u_{x0} + \frac{\partial N_1^{ip}}{\partial x} u_{x1} \right) A_x + \left(\frac{\partial N_0^{ip}}{\partial y} u_{x0} + \frac{\partial N_1^{ip}}{\partial y} u_{x1} \right) A_y \right. \\ &\quad \left. + \left(\frac{\partial N_0^{ip}}{\partial x} u_{y0} + \frac{\partial N_1^{ip}}{\partial x} u_{y1} \right) A_x + \left(\frac{\partial N_0^{ip}}{\partial y} u_{y0} + \frac{\partial N_1^{ip}}{\partial y} u_{y1} \right) A_y \right], \\ F_y &= -\Gamma_{ip} \left[\left(\frac{\partial N_0^{ip}}{\partial x} u_{y0} + \frac{\partial N_1^{ip}}{\partial x} u_{y1} \right) A_x + \left(\frac{\partial N_0^{ip}}{\partial y} u_{y0} + \frac{\partial N_1^{ip}}{\partial y} u_{y1} \right) A_y \right. \\ &\quad \left. + \left(\frac{\partial N_0^{ip}}{\partial x} u_{x0} + \frac{\partial N_1^{ip}}{\partial x} u_{x1} \right) A_x + \left(\frac{\partial N_0^{ip}}{\partial y} u_{x0} + \frac{\partial N_1^{ip}}{\partial y} u_{x1} \right) A_y \right]. \end{aligned}$$

For the edge-based diffusion operator, the value of ϕ is substituted for the component of velocity, u_i in the Eq. (3.43).

$$\frac{\partial u_i}{\partial x_{j\ ip}} = G_j^- u_i + [(u_{iR} - u_{iL}) - G_l^- u_i dx_l] \frac{A_j}{A_k dx_k}.$$

Common approaches in the cell-centered community are to use the projected nodal gradients for the $\frac{\partial u_j}{\partial x_i}$ stress component. However, in Nalu-Wind, the above form of equation is used.

Substituting the relations of the velocity gradients for the x and y-component of force above provides the following expression used for the viscous stress contribution:

$$\begin{aligned} F_x &= -\mu_{ip} \left[(G_x^- u_x A_x + G_y^- u_x A_y) + (u_{xR} - u_{xL}) \frac{A_x A_x + A_y A_y}{A_x dx + A_y dy} \right. \\ &\quad \left. - (G_x^- u_x dx + G_y^- u_x dy) \frac{A_x A_x + A_y A_y}{A_x dx + A_y dy} \right] \\ &\quad - \mu_{ip} \left[G_x^- u_x A_x + G_y^- u_y A_y + (u_{xR} - u_{xL}) \frac{A_x A_x}{A_x dx + A_y dy} \right. \\ &\quad \left. + (u_{yR} - u_{yL}) \frac{A_x A_y}{A_x dx + A_y dy} \right. \\ &\quad \left. - (G_x^- u_x dx + G_y^- u_x dy) \frac{A_x A_x}{A_x dx + A_y dy} \right. \\ &\quad \left. - (G_x^- u_y dx + G_y^- u_y dy) \frac{A_x A_y}{A_x dx + A_y dy} \right], \\ F_y &= -\mu_{ip} \left[(G_x^- u_y A_x + G_y^- u_y A_y) + (u_{yR} - u_{yL}) \frac{A_x A_x + A_y A_y}{A_x dx + A_y dy} \right. \\ &\quad \left. - (G_x^- u_y dx + G_y^- u_y dy) \frac{A_x A_x + A_y A_y}{A_x dx + A_y dy} \right] \\ &\quad - \mu_{ip} \left[G_y^- u_x A_x + G_y^- u_y A_y + (u_{yR} - u_{yL}) \frac{A_y A_y}{A_x dx + A_y dy} \right. \\ &\quad \left. + (u_{xR} - u_{xL}) \frac{A_y A_x}{A_x dx + A_y dy} \right. \\ &\quad \left. - (G_x^- u_y dx + G_y^- u_y dy) \frac{A_y A_y}{A_x dx + A_y dy} \right. \\ &\quad \left. - (G_x^- u_x dx + G_y^- u_x dy) \frac{A_y A_x}{A_x dx + A_y dy} \right], \end{aligned}$$

where above, the first \square and second \square represent the $\frac{\partial u_i}{\partial x_j} A_j$ and $\frac{\partial u_j}{\partial x_i} A_j$ contributions, respectively.

One can use this expression to recognize the ideal LHS sensitivities for row and columns for component u_i .

3.4 Advection Stabilization

In general, advection for both the edge and element-based scheme is identical with standard exception of the location of the integration points. The full advection term is simply written as,

$$ADV_\phi = \int \rho u_j \phi_{ip} A_j = \sum \dot{m} \phi_{ip}, \quad (3.44)$$

where ϕ is u_i , Z , h , etc.

The evaluation of ϕ_{ip} defines the advection stabilization choice. In general, the advection choice is a cell Peclet blending between higher order upwind (ϕ_{upw}) and a generalized un-stabilized central (Galerkin) operator, ϕ_{gcds} ,

$$\phi_{ip} = \eta \phi_{upw} + (1 - \eta) \phi_{gcds}. \quad (3.45)$$

In the above equation, η is a cell Peclet blending. The generalized central operator can take on a pure second order or pseudo fourth order form (see below). For the classic Peclet number functional form (see Equation (3.46)) a hybrid upwind factor, γ , can be used to ensure that no stabilization is added ($\eta = 0$) or that full upwind stabilization is included (as will be shown, even with limiter functions). The hybrid upwind factor allows one to modify the functional blending function; values of unity result in the normal blending function response in Figure Fig. 3.5; values of zero yield a pure central operator, i.e., blending function of zero; values \gg unity result in a blending function value of unity, i.e., pure upwind. The constant A is implemented with a value of 5. The value of this constant can not be changed via the input file. Note that this functional form is named the “classic” form within the input file.

The classic cell Peclet blending function is given by

$$\eta = \frac{\gamma \text{Pe}^2}{5 + \gamma \text{Pe}^2}. \quad (3.46)$$

The classic Peclet functional form makes it difficult to dial in the exact point at which the Peclet factor transitions from pure upwind to pure central. Therefore, an alternative form is provided that has a hyperbolic tangent functional form. This form allows one to specify the transition point and the width of the transition (see Equation (3.47)). The general tanh form is as follows,

$$\eta = \frac{1}{2} [(a + b) + (b - a) \tanh(\frac{\text{Pe} - c_{\text{trans}}}{c_{\text{width}}})] \quad (3.47)$$

Above, the constant c_{trans} represents the transition Peclet number while c_{width} represents the width of the transition. The value of λ is simply the shift between of the raw tanh function from zero while δ is the difference between the max Peclet factor (unity) and the minimum value prior to normalization. This approach ensures that the function starts at zero and asymptotes to unity,

$$\eta = \frac{1}{2} [1 + \tanh(\frac{\text{Pe} - c_{\text{trans}}}{c_{\text{width}}})].$$

The cell-Peclet number is computed for each sub-face in the element from the two adjacent left (L) and right (R) nodes,

$$\text{Pe} = \frac{\frac{1}{2} (u_{R,i} + u_{L,i}) (x_{R,i} - x_{L,i})}{\nu}.$$

A dot-product is implied by repeated indices.

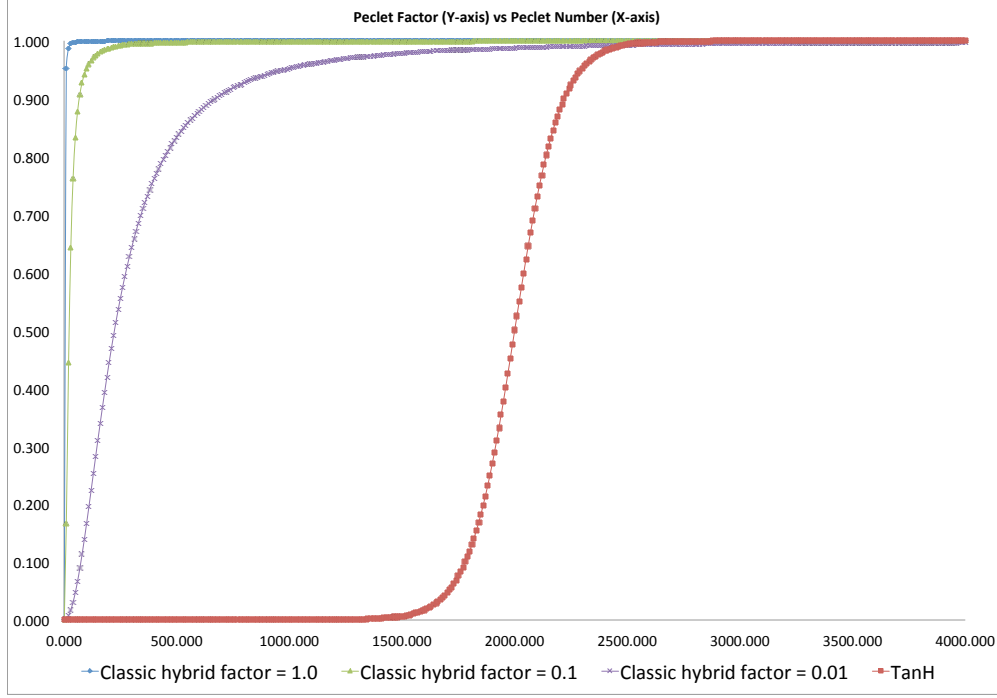


Fig. 3.5: Cell-Peclet number blending function outlining classic (varying the hybrid factor γ from 1.0, 0.1 and 0.01; again $A = 5$) and tanh functional form ($c_{trans} = 2000$ and $c_{width} = 200$).

The upwind operator, ϕ_{upw} is computed based on a blending of the extrapolated state (using the projected nodal gradient) and the linear interpolated state. Second or third order upwind is provided based on the value of α_{upw} blending

$$\begin{aligned}\phi_{upw} &= \alpha_{upw} \tilde{\phi}_{upw}^L + (1 - \alpha_{upw}) \phi_{cds}; \dot{m} > 0, \\ \alpha_{upw} \tilde{\phi}_{upw}^R + (1 - \alpha_{upw}) \phi_{cds}; \dot{m} < 0.\end{aligned}\quad (3.48)$$

The extrapolated value based on the upwinded left (ϕ^L) or right (ϕ^R) state,

$$\begin{aligned}\tilde{\phi}_{upw}^L &= \phi^L + d_j^L \frac{\partial \phi^L}{\partial x_j}, \\ \tilde{\phi}_{upw}^R &= \phi^R - d_j^R \frac{\partial \phi^R}{\partial x_j}.\end{aligned}\quad (3.49)$$

The distance vectors are defined based on the distances between the L/R points and the integration point (for both edge or element-based),

$$\begin{aligned}d_j^L &= x_j^{ip} - x_j^L, \\ d_j^R &= x_j^R - x_j^{ip}.\end{aligned}\quad (3.50)$$

In the case of all transported quantities, a Van Leer limiter of the extrapolated value is supported and can be activated within the input file (using the solution options “limiter” specification).

Second order central is simply written as a linear combination of the nodal values,

$$\phi_{cds} = \sum N_k^{ip} \phi_k. \quad (3.51)$$

where N_k^{ip} is either evaluated at the subcontrol surface or edge midpoint. In the case of the edge-based scheme, the edge midpoint evaluation provides for a skew symmetric form of the operator.

The generalized central difference operator is provided by blending with the extrapolated values and second order Galerkin,

$$\phi_{gcds} = \frac{1}{2} \left(\hat{\phi}_{upw}^L + \hat{\phi}_{upw}^R \right), \quad (3.52)$$

where,

$$\begin{aligned} \hat{\phi}_{upw}^L &= \alpha \tilde{\phi}_{upw}^L + (1 - \alpha) \phi_{cds}, \\ \hat{\phi}_{upw}^R &= \alpha \tilde{\phi}_{upw}^R + (1 - \alpha) \phi_{cds}. \end{aligned} \quad (3.53)$$

The value of α provides the type of psuedo fourth order stencil and is specified in the user input file.

The above set of advection operators can be used to define an idealized one dimensional stencil denoted by $(i - 2, i - 1, i, i + 1, i + 2)$, where i represents the particular row for the given transported variable. Below, in the table, the stencil can be noted for each value of α and α_{upw} .

$i - 2$	$i - 1$	i	$i + 1$	$i + 2$	α	α_{upw}
0	$-\frac{1}{2}$	0	$+\frac{1}{2}$	0	0	n/a
$+\frac{1}{8}$	$-\frac{6}{8}$	0	$+\frac{6}{8}$	$-\frac{1}{8}$	$\frac{1}{2}$	n/a
$+\frac{1}{12}$	$-\frac{8}{12}$	0	$+\frac{8}{12}$	$-\frac{1}{12}$	$\frac{2}{3}$	n/a
$+\frac{1}{4}$	$-\frac{5}{4}$	$+\frac{3}{4}$	$+\frac{1}{4}$	0	$\dot{m} > 0$	1
0	$-\frac{1}{4}$	$-\frac{3}{4}$	$+\frac{5}{4}$	$-\frac{1}{4}$	$\dot{m} < 0$	1
$+\frac{1}{6}$	$-\frac{6}{6}$	$+\frac{3}{6}$	$+\frac{2}{6}$	0	$\dot{m} > 0$	$\frac{1}{2}$
0	$-\frac{2}{6}$	$-\frac{3}{6}$	$+\frac{6}{6}$	$-\frac{1}{6}$	$\dot{m} < 0$	$\frac{1}{2}$

It is noted that by varying these numerical parameters, both high quality, low dissipation operators suitable for LES usage or limited, monotonic operators suitable for RANS modeling can be accomodated.

3.5 Pressure Stabilization

A number of papers describing the pressure stabilization approach that Nalu-Wind uses are in the open literature, Domino, [Dom06], [Dom08], [Dom14]. Nalu-Wind supports an incremental fourth order approximate projection scheme with time step scaling. By scaling, it is implied that a time scale based on either the physical time step or a combined elemental advection and diffusion time scale based on element length along with advection and diffusional parameters. An alternative to the appproximate projection concept is to view the method as a variational multiscale (VMS) method whereby the momentum residual augments the continuity equation. This allows for a diagonal entry for the pressure degree of freedom.

Here, the fine-scale momentum residual is written in terms of a projected momentum residual evaluated at the Gauss point,

$$\mathbf{R}(u_i) = \left(\frac{\partial p}{\partial x_j} - G_j p \right). \quad (3.54)$$

The above equation is derived simply by writing a fine-scale momentum equation at the Gauss-points and using the nodal projected residual to reconstruct the individual terms. Therefore, the continuity equation solved, using the VMS-based projected momentum residual, is

$$\int \frac{\partial \bar{p}}{\partial t} dV + \int (\bar{\rho} \hat{u}_i + \tau G_i \bar{P}) n_i dS = \int \tau \frac{\partial \bar{P}}{\partial x_i} n_i dS.$$

Above, $G_i \bar{P}$ is defined as a L2 nodal projection of the pressure gradient. Note that the notion of a provisional velocity, \hat{u}_i , is used to signify that this velocity is the product of the momentum solve. The difference between the projected

nodal gradient interpolated to the gauss point and the local gauss point pressure gradient provides a fourth order pressure stabilization term. This term can also be viewed as an algebraic form for the momentum residual. For the continuity equation only, a series of element-based options that shift the integration points to the edges of the iterated element is an option.

3.5.1 The Role of \dot{m}

In all of the above equations, the advection term is written in terms of a linearized mass flow rate including a sum over all subcontrol surface integration points, Eq (3.44). The mass flow rate includes the full set of stabilization terms obtained from the continuity solve,

$$\dot{m} = \left(\bar{\rho} \hat{u}_i + \tau G_i \bar{P} - \tau \frac{\partial \bar{P}}{\partial x_i} \right) n_i dS.$$

The inclusion of the pressure stabilization terms in the advective transport for the primitives is a required step for ensuring that the advection velocity is mass conserving. In practice, the mass flow rate is stored at each integration point in the mesh (edge midpoints for the edge-based scheme and subcontrol surfaces for the element-based scheme). When the mixed CVFEM/EBVC scheme is used, the continuity equation solves for a subcontrol-surface value of the mass flow rate. These values are assembled to the edge for use in the EBVC discretization approach. Therefore, the storage for mass flow rate is higher.

3.6 RTE Stabilization

The RTE is solved using the method of discrete ordinates using the symmetric Thurgood quadrature set. The discrete ordinates method is one in which discrete directions of the intensity are solved. The quadrature order, N , defines the number of ordinate directions that are solved in a given iteration. In the case of non-scattering media, this results in a set of decoupled linear partial differential equations. For the symmetric Thurgood set, the number of ordinate directions is given by $8N^2$. Values of N that are required for suitable accuracy starts at $N = 2$ with more than adequate resolution at $N = 4$.

For each ordinate direction, a weight is provided, w_k (not to be confused with the test function w). For each intensity ordinate direction, I_k , integrated quantities such as scalar flux and radiative heat flux are computed as,

$$G = \sum I_k w_k$$

and,

$$q_j = \sum I_k s_j w_k.$$

The stabilization that is used in the RTE equation can be placed in the class of residual-based stabilization. In this particular implementation, the scaled residual of the RTE equation is added. This implementation has its roots in the classic variational multiscale (VMS).

In the VMS framework, the degree of freedom is decomposed in terms of its resolved and fine scale, $I + I'$. Without specific definition of the test function, the weighted residual statement for the RTE within a VMS framework is given by,

$$\int w \left(s_i \frac{\partial}{\partial x_i} (I(s) + I'(s)) + (\mu_a + \mu_s) (I(s) + I'(s)) - \frac{\mu_a \sigma T^4}{\pi} - \frac{\mu_s}{4\pi} G \right) dV = 0. \quad (3.55)$$

Grouping resolved and fine scale terms results in an equation takes the form of a standard Galerkin contribution in addition to the fine structure statement,

$$\begin{aligned} \int w \left(s_i \frac{\partial}{\partial x_i} I(s) + (\mu_a + \mu_s) I - \frac{\mu_a \sigma T^4}{\pi} - \frac{\mu_s}{4\pi} G \right) dV \\ + \int w \left(s_i \frac{\partial}{\partial x_i} I'(s) + (\mu_a + \mu_s) I' \right) dV = 0. \end{aligned} \quad (3.56)$$

Note that the isotropic source term has not contributed to the VMS framework other than through the right hand source term.

In general, gradients in the fine scale quantity are to be avoided. Therefore, the first term in the second line of Eq. (3.56) is integrated by parts to yield the following form (note the boundary term, \int_{Γ} that is shown below is frequently dropped)

$$\begin{aligned} & \int w \left(s_i \frac{\partial}{\partial x_i} I(s) + (\mu_a + \mu_s) I - \frac{\mu_a \sigma T^4}{\pi} - \frac{\mu_s}{4\pi} G \right) dV \\ & - \int I' s_i \frac{\partial w}{\partial x_i} dV + \int_{\Gamma} w s_i I' n_i dS + \int w (\mu_a + \mu_s) I' dV = 0. \end{aligned} \quad (3.57)$$

The following ansatz, which now includes the classic stabilization parameter, τ , provides closure of the above fine scale equation,

$$I' = -\tau \left(s_i \frac{\partial}{\partial x_i} I(s) + (\mu_a + \mu_s) I(s) - \frac{\mu_a \sigma T^4}{\pi} - \frac{\mu_s}{4\pi} G \right) = -\tau R(s) \quad (3.58)$$

Substituting Eq. (3.58) into Eq. (3.57) yields,

$$\begin{aligned} & \int w \left(s_i \frac{\partial}{\partial x_i} I(s) + (\mu_a + \mu_s) I - \frac{\mu_a \sigma T^4}{\pi} - \frac{\mu_s}{4\pi} G \right) dV \\ & + \int \tau s_i \frac{\partial w}{\partial x_i} R(s) dV - \int_{\Gamma} \tau w R(s) s_i n_i dS - \int \tau w (\mu_a + \mu_s) R(s) dV = 0. \end{aligned} \quad (3.59)$$

In the above equation, the residual of the intensity equation for ordinate s is denoted by $R(s)$. A compact form of the equation is provided by defining a modified test function, \tilde{w} , (again note retention of the stabilized boundary term)

$$\begin{aligned} & \int \tilde{w} \left(s_i \frac{\partial}{\partial x_i} I(s) + (\mu_a + \mu_s) I - \frac{\mu_a \sigma T^4}{\pi} - \frac{\mu_s}{4\pi} G \right) dV \\ & - \beta \int_{\Gamma} \tau w R(s) s_i n_i dS = 0. \end{aligned} \quad (3.60)$$

where \tilde{w} is simply equal to,

$$\tilde{w} = w + \tau \left(s_j \frac{\partial w}{\partial x_j} + \alpha (\mu_a + \mu_s) w \right). \quad (3.61)$$

When $\alpha = -1$, we have the above VMS derivation; for $\alpha = 1$, Galerkin Least Squares is realized; finally for $\alpha = 0$, we have SUPG. For any formulation other than VMS, the residual contribution at the boundaries of the domain is dropped ($\beta = 0$).

The full residual-based equation is placed in divergence form,

$$\begin{aligned} & \int \tilde{w} \left(\frac{\partial}{\partial x_i} s_i I(s) + (\mu_a + \mu_s) I(s) - \frac{\mu_a \sigma T^4}{\pi} - \frac{\mu_s}{4\pi} G \right) dV \\ & - \beta \int_{\Gamma} \tau w R(s) s_i n_i dS = 0. \end{aligned} \quad (3.62)$$

and split into its Galerkin and stabilized contributions,

$$\begin{aligned} & \int w \left(\frac{\partial}{\partial x_i} s_i I(s) + (\mu_a + \mu_s) I(s) - \frac{\mu_a \sigma T^4}{\pi} - \frac{\mu_s}{4\pi} G \right) dV \\ & + \int \tau s_j \frac{\partial w}{\partial x_j} R(s) dV \\ & + \alpha \int \tau w (\mu_a + \mu_s) R(s) dV \\ & - \beta \int_{\Gamma} \tau w R(s) s_i n_i dS = 0. \end{aligned} \quad (3.63)$$

Note that the first term in the above equation is integrated by parts,

$$\int w \frac{\partial}{\partial x_i} s_i I(s) dV = - \int I(s) s_i \frac{\partial w}{\partial x_i} dV + \int_{\Gamma} w s_i I(s) n_i dS.$$

Again, the usage of Γ provides emphasis that the contribution is a boundary (exposed face) condition. Therefore, the full VMS-based stabilized RTE equation is as follows,

$$\begin{aligned} & \int \left(-I(s) s_i \frac{\partial w}{\partial x_i} + (\mu_a + \mu_s) I(s) - \frac{\mu_a \sigma T^4}{\pi} - \frac{\mu_s}{4\pi} G \right) dV \\ & + \int_{\Gamma} w s_i I(s) n_i dS \\ & + \int \tau s_j \frac{\partial w}{\partial x_j} R(s) dV \\ & + \alpha \int \tau w (\mu_a + \mu_s) R(s) dV \\ & - \beta \int_{\Gamma} \tau w R(s) s_i n_i dS = 0. \end{aligned} \quad (3.64)$$

3.6.1 Definition of the test function

Following the work of Martinez, [Mar05], the test function is chosen to be a piecewise-constant value within the control volume, $w = w_I$ and zero outside of this control volume (Heaviside). A key property of this function, as pointed out by Martinez, is that its gradient is a distribution of delta functions on the control volume boundary:

$$\frac{\partial w_I}{\partial x_i} = -\mathbf{n}_I \delta(\mathbf{x} - \mathbf{x}_{\Gamma_I}) \quad (3.65)$$

where Γ_I is boundary of control volume I and \mathbf{n}_I is the outward normal on that boundary. Substituting this relationship into the residual equation provides the final form of vertex-centered finite volume RTE stabilized equation,

$$\begin{aligned} & \int I(s) s_i n_i dS + \int \left((\mu_a + \mu_s) I(s) - \frac{\mu_a \sigma T^4}{\pi} - \frac{\mu_s}{4\pi} G \right) dV \\ & + \int_{\Gamma} s_i I(s) n_i dS \\ & - \int \tau R(s) s_i n_i dS + \alpha \int \tau (\mu_a + \mu_s) R(s) dV - \beta \int_{\Gamma} \tau R(s) s_i n_i dS = 0. \end{aligned} \quad (3.66)$$

Given this equation, either an edge-based or element-based scheme can be used. For $\alpha = 0$ and $\beta = 0$, it is noted that classic SUCV is obtained. The second line of Eq. (3.66) represents a boundary contribution. This is where the intensity boundary condition (Eq. (3.126)) is applied. As noted in the RTE equation section, when $s_j n_j$ is greater than zero, the interpolated intensity based on the surface nodal values is used. However, when $s_j n_j$ is less than zero, the intensity boundary condition value is used. Since the original RTE equation was integrated by parts, a natural surface flux contribution is applied. In alternative discretization approaches, e.g., the SUPG FEM-based Sierra Thermal Radiation Module: Syrinx code, the RTE is not integrated by parts. Therefore, no boundary term exists, and, therefore, a dirichlet bc is applied. At corner nodes, this approach can lead to non-intuitive approaches since the corner node might have surface facets that are both incoming and outgoing. Weak integration of the flux term eliminated this complexity.

3.6.2 The form of τ

The value of the stabilization parameter τ can take on a variety of forms. A classic derivation provides the form of τ to be broken out into two forms, $\tau_{adv} = \frac{h}{2}$ and $\tau_{diff} = \frac{1}{(\mu_a + \mu_s)}$. An ad-hoc blending is given by,

$$\tau = \left(\frac{1}{\frac{2}{h^2} + (\mu_a + \mu_s)^2} \right)^{\frac{1}{2}} \quad (3.67)$$

Finally, the classic GFEM form of τ is given by use of the metric tensor for the element mapping is noted,

$$\tau = \beta^* [s_i g_{ij} s_j]^{-\frac{1}{2}}, \quad (3.68)$$

with β^* equal to unity for SUCV and $\frac{2}{15^{\frac{1}{2}}}$ for FEM.

3.6.3 Pure Edge-based Upwind Method

The residual-based stabilization approach can lead to predicting negative intensities. This is simply due to the fact that the stabilization approach (SUPG) is a linear approach. Extensions of this residual-based stabilization to include a discontinuity capturing operator (DCO) are underway. This adds a non-linear stabilization approach that will, hopefully, eliminate negative intensity predictions.

Alternatively, a first order upwind approach has been implemented by using EBVC discretization. At this point, no higher order upwind extensions have been implemented. For the upwind implementation, the equation solved is,

$$\begin{aligned} \int I(s) s_i n_i dS + \int \left((\mu_a + \mu_s) I(s) - \frac{\mu_a \sigma T^4}{\pi} - \frac{\mu_s}{4\pi} G \right) dV \\ + \int_{\Gamma} s_i I(s) n_i dS = 0. \end{aligned} \quad (3.69)$$

In the above equation, the “advection operator”, $I(s) s_i n_i dS$ is approximated as using the “upwind” intensity, e.g., if $s_j n_j$ is greater than zero, the left nodal value is used.

3.6.4 Finite Element SUPG Form

For the FEM, the test function is the standard weighting. Assuming a pure SUPG formulation, i.e., $\alpha = \beta = 0$ in Equation (3.64), thereby reducing the final form to the following:

$$\begin{aligned} \int \left(-I(s) s_i \frac{\partial w}{\partial x_i} + w \left[(\mu_a + \mu_s) I(s) - \frac{\mu_a \sigma T^4}{\pi} - \frac{\mu_s}{4\pi} G \right] \right) dV \\ + \int_{\Gamma} w s_i I(s) n_i dS \\ + \int \tau s_j \frac{\partial w}{\partial x_j} R(s) dV \end{aligned} \quad (3.70)$$

The weak boundary condition is applied in a similar manner as with the CVFEM and EBVC form, however, using the appropriate FEM test function definition. Finally, the form of τ follows the above CVFEM form.

3.7 Nonlinear Stabilization Operator (NSO)

An alternative to classic Peclet number blending is the usage of a discontinuity capturing operator (DCO), or in the low Mach context a nonlinear stabilization operator (NSO). In this method, an artificial viscosity is defined that is a function of the local residual and scaled computational gradients. Viable usages for the NSO can be advection/diffusion problems in addition to the aforementioned RTE VMS approach.

The formal finite element kernel for a NSO approach is as follows,

$$\sum_e \int_{\Omega} \nu(\mathbf{R}) \frac{\partial w}{\partial x_i} g^{ij} \frac{\partial \phi}{\partial x_j} d\Omega, \quad (3.71)$$

where $\nu(\mathbf{R})$ is the artificial viscosity which is a function of the pde fine-scale residual and g^{ij} is the covariant metric tensor).

For completeness, the covariant and contravariant metric tensor are given by,

$$g^{ij} = \frac{\partial x_i}{\partial \xi_k} \frac{\partial x_j}{\partial \xi_k}, \quad (3.72)$$

and

$$g_{ij} = \frac{\partial \xi_k}{\partial x_i} \frac{\partial \xi_k}{\partial x_j}, \quad (3.73)$$

where $\xi = (\xi_1, \xi_2, \xi_3)^T$. The form of $\nu(\mathbf{R})$ currently used is as follows,

$$\nu = \sqrt{\frac{\mathbf{R}_k \mathbf{R}_k}{\frac{\partial \phi}{\partial x_i} g^{ij} \frac{\partial \phi}{\partial x_j}}}. \quad (3.74)$$

The classic paper by Shakib ([SHZ91]) represents the genesis of this method which was done in the accoustically compressible context.

A residual for a classic advection/diffusion/source pde is simply the fine scale residual computed at the gauss point,

$$\hat{\mathbf{R}} = \frac{\partial \rho \phi}{\partial t} + \frac{\partial}{\partial x_j} (\rho u_j \phi - \mu^{eff} \frac{\partial \phi}{\partial x_j}) - S \quad (3.75)$$

Note that the above equation requires a second derivative whose source is the diffusion term. The first derivative is generally determined by using projected nodal gradients. As will be noted in the pressure stabilization section, the advection term carries the pressure stabilization terms. However, in the above equation, these terms are not explicitly noted. Therefore, an option is to subtract the fine scale continuity equation to obtain the final general form of the source term,

$$\mathbf{R} = \hat{\mathbf{R}} - \phi \left(\frac{\partial \rho}{\partial t} + \frac{\partial \rho u_j}{\partial x_j} \right). \quad (3.76)$$

An alternative to the fine-scale PDE is a form that is found by differencing the linearized form of the residual from the nonlinear residual,

$$\mathbf{R} = \frac{\partial \rho u_j \phi}{\partial x_j} - \left(\phi \frac{\partial \rho u_j}{\partial x_j} + \rho u_j \frac{\partial \phi}{\partial x_j} \right). \quad (3.77)$$

The above resembles a commutation error in the nonlinear advection term.

In general, the NSO- ν is prone to percision issues when the gradients are very close to zero. As such, the value of ν is limited to a first-order like value. This parameter is proposed as follows: if an operator were written as a Galerkin (un-stabilized) plus a diffusion operator, what is the value of the diffusion coefficient such that first-order upwind is obtained for the combined operator? This upwind limited value of ν provides the highest value that this coefficient can (or should) be. The current form of the limited upwind ν is as follows,

$$\nu^{upw} = C_{upw} (\rho u_i g_{ij} \rho u_j)^{\frac{1}{2}} \quad (3.78)$$

where C_{upw} is generally taked to be 0.1.

Using a piecewise-constant test function suitable for CVFEM and EBVC schemes (the reader is refered to the VMS RTE section), Eq. (3.71) can be written as,

$$- \sum_e \int_{\Gamma} \nu(\mathbf{R}) g^{ij} \frac{\partial \phi}{\partial x_j} n_i dS. \quad (3.79)$$

A fourth order form, which writes the stabilization as the difference between the Gauss-point gradient and the projected nodal gradient interpolated to the Gauss-point, is also supported,

$$- \sum_e \int_{\Gamma} \nu(\mathbf{R}) g^{ij} \left(\frac{\partial \phi}{\partial x_j} - G_j \phi \right) n_i dS. \quad (3.80)$$

3.7.1 NSO Based on Kinetic Energy Residual

An alternative formulation explored is to share the general kernel form shown in Equation (3.80), however, compute ν based on a fine-scale kinetic energy residual. In this formulation, the fine-scale kinetic energy residual is obtained from the fine-scale momentum residual dotted with velocity. As with the continuity stabilization approach, the fine-scale momentum residual is provided by Equation (3.81). Therefore, the fine-scale kinetic energy is written as,

$$\mathbf{R}_{ke} = \frac{u_j (\frac{\partial p}{\partial x_j} - G_j p)}{2}, \quad (3.81)$$

while the denominator for ν now includes the gradient in ke,

$$\nu = \sqrt{\frac{\mathbf{R}_{ke} \mathbf{R}_{ke}}{\frac{\partial ke}{\partial x_i} g^{ij} \frac{\partial ke}{\partial x_j}}}. \quad (3.82)$$

The kinetic energy is simply given by,

$$ke = \frac{u_k u_k}{2} \quad (3.83)$$

The kinetic energy form of ν is used for all equation sets with transformation by usage of a turbulent Schmidt/Prandtl number.

3.7.2 Local or Projected NSO Diffusive Flux Coefficient

While the NSO kernel is certainly evaluated at the subcontrol surfaces, the evaluation of ν can be computed by a multitude of approaches. For example, the artificial diffusive flux coefficient can be computed locally (with local residuals and local metric tensors) or can be projected to the nodes (via an L_{oo} or L_2 projection) and re-interpolated to the gauss points. The former results in a sharper local value while the later results in a more filtered-like value. The code base only supports a local NSO ν calculation.

3.7.3 General Findings

In general, the NSO approach seems to work best when running the fourth-order option as the second-order implementation still looks more diffuse. When compared to the standard MUSCL-limited scheme, the NSO is the preferred choice. More work is underway to improve stabilization methods. Note that a limited set of NSOs are activated in the code base with specific interest on scalar transport, e.g., momentum, mixture fraction and static enthalpy transport. When using the 4th order method, the consistent mass matrix approach for the projected nodal gradients is supported for higher order.

3.7.4 NSO as a Turbulence Model

The kinetic energy residual form has been suggested to be used as a turbulence model (Guermond and Larios, 2015). However, inspection of the above NSO kernel form suggests that the model form is not symmetric. Rather, in the context of turbulence modeling, is closer to the metric tensor acting on the difference between the rate of strain and antisymmetric tensor. As such, the theory developed, e.g., for eigenvalue perturbations of the stress tensor (see Jofre and Domino, 2017) can not be applied. In this section, a new form of the NSO is provided in an effort to be used for an LES closure.

In this proposed NSO formulation, the subgrid stress tensor, $\tau_{ij}^{sgs} = \overline{u_i u_j} - \bar{u}_i \bar{u}_j$, is given by,

$$\tau_{ij}^{sgs} = -2\rho\nu g^{ij} (S_{ij} - \frac{1}{3} \frac{\partial u_k}{\partial x_k} \delta_{ij}) = -2\rho\nu g^{ij} S_{ij}^*. \quad (3.84)$$

Interestingly, the units of ν are of an inverse time scale while the product $2\rho\nu g^{ij}$ can be viewed as a non-isotropic eddy viscosity, μ_{ij}^t .

The first order clipping may be relaxed by defining ν as,

$$\nu = \frac{|\mathbf{R}_{ke}|}{||ke||_\infty}. \quad (3.85)$$

The above form would be closer to what Guermond uses and would avoid the divide-by-zero noted in regions of uniform flow.

3.8 Turbulence Modeling

Unlike a RANS approach which models most or all of the turbulent fluctuations, LES directly solves for all resolved turbulent length scales and only models the smallest scales below the grid size. In this way, a majority of the problem-dependent, energy-containing turbulent structure is directly solved in a model-free fashion. The subgrid scales are closer to being isotropic than the resolved scales, and they generally act to dissipate turbulent kinetic energy cascaded down from the larger scales in momentum-driven turbulent flows. Modeling of these small scales is generally more straightforward than RANS approaches, and overall solutions are usually more tolerant to LES modeling errors because the subgrid scales comprise such a small portion of the overall turbulent structure. While LES is generally accepted to be much more accurate than RANS approaches for complex turbulent flows, it is also significantly more expensive than equivalent RANS simulations due to the finer grid resolution required. Additionally, since LES results in a full unsteady solution, the simulation must be run for a long time to gather any desired time-averaged statistics. The tradeoff between accuracy and cost must be weighed before choosing one method over the other.

The separation of turbulent length scales required for LES is obtained by using a spatial filter rather than the RANS temporal filter. This filter has the mathematical form

$$\overline{\phi(\mathbf{x}, t)} \equiv \int_{-\infty}^{+\infty} \phi(\mathbf{x}', t) G(\mathbf{x}' - \mathbf{x}) d\mathbf{x}', \quad (3.86)$$

which is a convolution integral over physical space \mathbf{x} with the spatially-varying filter function G . The filter function has the normalization property $\int_{-\infty}^{+\infty} G(\mathbf{x}) d\mathbf{x} = 1$, and it has a characteristic length scale Δ so that it filters out turbulent length scales smaller than this size. In the present formulation, a simple “box filter” is used for the filter function,

$$G(\mathbf{x}' - \mathbf{x}) = \begin{cases} 1/V & : (\mathbf{x}' - \mathbf{x}) \in \mathcal{V} \\ 0 & : \text{otherwise} \end{cases},$$

where V is the volume of control volume \mathcal{V} whose central node is located at \mathbf{x} . This is essentially an unweighted average over the control volume. The length scale of this filter is approximated by $\Delta = V^{1/3}$. This is typically called the grid filter, as it filters out scales smaller than the computational grid size.

Similar to the RANS temporal filter, a variable can be represented in terms of its filtered and subgrid fluctuating components as

$$\phi = \bar{\phi} + \phi'.$$

For most forms of the filter function $G(\mathbf{x})$, repeated applications of the grid filter to a variable do not yield the same result. In other words, $\bar{\bar{\phi}} \neq \bar{\phi}$ and therefore $\bar{\phi}' \neq 0$, unlike with the RANS temporal averages.

As with the RANS formulation, modeling is much simplified in the presence of large density variations if a Favre-filtered approach is used. A Favre-filtered variable $\tilde{\phi}$ is defined as

$$\tilde{\phi} \equiv \frac{\overline{\rho\phi}}{\bar{\rho}}$$

and a variable can be decomposed in terms of its Favre-filtered and subgrid fluctuating component as

$$\phi = \tilde{\phi} + \phi''.$$

Again, note that the useful identities for the Favre-filtered RANS variables do not apply, so that $\tilde{\tilde{\phi}} \neq \tilde{\phi}$ and $\overline{\phi''} \neq 0$. The Favre-filtered approach is used for all LES models in Nalu-Wind.

3.8.1 Standard Smagorinsky LES Model

The standard Smagorinsky LES closure model approximates the subgrid turbulent eddy viscosity using a mixing length-type model, where the LES grid filter size Δ provides a natural length scale. The subgrid eddy viscosity is modeled simply as (Smagorinsky)

$$\mu_t = \rho (C_s \Delta)^2 |\tilde{S}|, \quad (3.87)$$

The constant coefficient C_s typically varies between 0.1 and 0.24 and should be carefully tuned to match the problem being solved (Rogallo and Moin, [RM84]). The default value of 0.17 is assigned in the code base.

Although this model is desirable due to its simplicity and efficiency, care should be taken in its application. It is known to predict subgrid turbulent eddy viscosity proportional to the shear rate in the flow, independent of the local turbulence intensity. Non-zero subgrid turbulent eddy viscosity is even predicted in completely laminar regions of the flow, sometimes even preventing a natural transition to turbulence. The model also does not asymptotically replicate near wall behavior without either dampening or a dynamic procedure.

3.8.2 Wall Adapting Local Eddy-Viscosity, WALE

The WALE model of Ducros et al., [DNP98], properly captures the asymptotic behavior for flows that are wall bounded. In this model, the turbulent viscosity is given by,

$$\mu_t = \rho (C_w \Delta)^2 \frac{(S_{ij}^d S_{ij}^d)^{3/2}}{(S_{ij} S_{ij})^{5/2} + (S_{ij}^d S_{ij}^d)^{5/4}}, \quad (3.88)$$

with the constant C_w of 0.325 and a standard filter, Δ related to the volume, $V^{1/3}$. The rate of strain tensor is defined as,

$$S_{ij} = \frac{1}{2} \left(\frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i} \right) \quad (3.89)$$

while S_{ij}^d is,

$$S_{ij}^d = \frac{1}{2} (g_{ij}^2 + g_{ji}^2) - \frac{1}{3} \delta_{ij} g_{kk}^2. \quad (3.90)$$

Finally, the velocity gradient squared terms are

$$g_{ij}^2 = \frac{\partial u_i}{\partial x_k} \frac{\partial u_k}{\partial x_j} \quad (3.91)$$

and

$$g_{ji}^2 = \frac{\partial u_j}{\partial x_k} \frac{\partial u_k}{\partial x_i}. \quad (3.92)$$

3.8.3 One Equation k^{sgs}

See k^{sgs} PDE section.

3.8.4 SST RANS Model

As noted, Nalu-Wind does support a SST RANS-based model (the reader is referred to the SST equation set description).

3.8.5 Hybrid RANS/LES Models

Nalu-Wind supports the Active Model Split (AMS) hybrid RANS/LES turbulence model [HOM20]. The reader is referred to the AMS equation set for more details.

3.8.6 Wall Models

Flows are either expected to be fully resolved or, alternatively, under-resolved where wall functions are used. A classic law of the wall has been implemented in Nalu-Wind. Wall models to handle adverse pressure gradients are planned. For more information of the form of wall models, please refer to the *boundary condition* section of this manual.

3.9 Supported Boundary Conditions

3.9.1 Inflow Boundary Condition

Continuity

Continuity uses a flux boundary condition with the incoming mass flow rate based on the user specified values for velocity,

$$\dot{m}_c = \rho^{spec} u_j^{spec} A_j.$$

As this is a vertex-based code, at inflow and Dirichlet wall boundary locations, the continuity equation uses the specified velocity within the inflow boundary condition block.

Momentum, Mixture Fraction, Enthalpy, Species, k_{sgs} , \mathbf{k} and ω

These degree-of-freedom (DOFs) each use a Dirichlet value with the specified user value. For all Dirichlet values, the row is zeroed with a unity placed on the diagonal. The residual is zeroed and set to the difference between the current value and user specified value.

3.9.2 Wall Boundary Conditions

Continuity

Continuity uses a no-op.

Momentum

When resolving the boundary layer, Momentum again uses a no-slip Dirichlet condition., e.g., $u_i = 0$.

In the case of a wall model, a classic wall function is applied. The wall shear stress enters the discretization of the momentum equations by the term

$$\int \tau_{ij} n_j dS = -F_{wi}. \quad (3.93)$$

Wall functions are used to prescribe the value of the wall shear stress rather than resolving the boundary layer within the near-wall domain. The fundamental momentum law of the wall formulation, assuming fully-developed turbulent flow near a no-slip wall, can be written as,

$$u^+ = \frac{u_{||}}{u_\tau} = \frac{1}{\kappa} \ln(Ey^+), \quad (3.94)$$

where u^+ is defined by the the near-wall parallel velocity, $u_{||}$, normalized by the wall friction velocity, u_τ . The wall friction velocity is related to the turbulent kinetic energy by,

$$u_\tau = C_\mu^{1/4} k^{1/2}. \quad (3.95)$$

by assuming that the production and dissipation of turbulence is in local equilibrium. The wall friction velocity is also computed given the density and wall shear stress,

$$u_\tau = \left(\frac{\tau_w}{\rho}\right)^{0.5}.$$

The normalized perpendicular distance from the point in question to the wall, y^+ , is defined as the following:

$$y^+ = \frac{\rho Y_p}{\mu} \left(\frac{\tau_w}{\rho}\right)^{1/2} = \frac{\rho Y_p u_\tau}{\mu}. \quad (3.96)$$

The classical law of the wall is as follows:

$$u^+ = \frac{1}{\kappa} \ln(y^+) + C, \quad (3.97)$$

where κ is the von Karman constant and C is the dimensionless integration constant that varies based on authorship and surface roughness. The above expression can be re-written as,

$$u^+ = \frac{1}{\kappa} \ln(y^+) + \frac{1}{\kappa} \ln(\exp(\kappa C)), \quad (3.98)$$

or simplified to the following expression:

$$\begin{aligned} u^+ &= \frac{1}{\kappa} (\ln(y^+) + \ln(\exp(\kappa C))) \\ &= \frac{1}{\kappa} \ln(Ey^+). \end{aligned} \quad (3.99)$$

In the above equation, E , is referred to in the text as the dimensionless wall roughness parameter and is described by,

$$E = \exp(\kappa C). \quad (3.100)$$

In Nalu-Wind, κ is set to the value of 0.42 while the value of E is set to 9.8 for smooth walls (White suggests values of $\kappa = 0.41$ and $E = 7.768$). The viscous sublayer is assumed to extend to a value of $y^+ = 11.63$.

The wall shear stress, τ_w , can be expressed as,

$$\tau_w = \rho u_\tau^2 = \rho u_\tau \frac{u_{||}}{u^+} = \frac{\rho \kappa u_\tau}{\ln(Ey^+)} u_{||} = \lambda_w u_{||}, \quad (3.101)$$

where λ_w is simply the grouping of the factors from the law of the wall. For values of y^+ less than 11.63, the wall shear stress is given by,

$$\tau_w = \mu \frac{u_{||}}{Y_p}. \quad (3.102)$$

The force imparted by the wall, for the i_{th} component of velocity, can be written as,

$$F_{w,i} = -\lambda_w A_w u_{i||}, \quad (3.103)$$

where A_w is the total area over which the shear stress acts.

The use of a general, non-orthogonal mesh adds a slight complexity to specifying the force imparted on the fluid by the wall. As shown in Equation (3.103), the velocity component parallel to the wall must be determined. Use of the unit normal vector, n_j , provides an easy way to determine the parallel velocity component by the following standard vector projection:

$$\Pi_{ij} = [\delta_{ij} - n_i n_j]. \quad (3.104)$$

Carrying out the projection of a general velocity, which is not necessarily parallel to the wall, yields the velocity vector parallel to the wall,

$$u_{i||} = \Pi_{ij} u_j = u_i (1 - n_i^2) - \sum_{j=1; j \neq i}^n u_j n_i n_j. \quad (3.105)$$

Note that the component that acts on the particular i^{th} component of velocity,

$$-\lambda_w A_w (1 - n_i n_i) u_{i||}, \quad (3.106)$$

provides a form that can be potentially treated implicitly; i.e., in a way to augment the diagonal dominance of the central coefficient of the i^{th} component of velocity. The use of residual form adds a slight complexity to this implicit formulation only in that appropriate right-hand-side source terms must be added.

Mixture Fraction

If a value is specified for each quantity within the wall boundary condition block, a Dirichlet condition is applied. If no values are specified, a zero flux condition is applied.

Enthalpy

If the temperature is specified within the wall boundary condition block, a Dirichlet condition is always specified. Wall functions for enthalpy transport have not yet been implemented.

The simulation tool supports multi-physics coupling via conjugate heat transfer and radiative heat transfer. Coupling parameters required for the thermal boundary condition are post processed by the fluids or PMR Realm. For conjugate and radiative coupling, the thermal solve provides the surface temperature. From the surface temperature, a wall enthalpy is computed and used.

Thermal Heat Conduction

If a temperature is specified in the wall block, and the surface is not an interface condition, then a Dirichlet approach is used. If conjugate heat transfer is included, then the boundary condition applied is as follows,

$$-\kappa \frac{\partial T}{\partial x_j} n_j dS = h(T - T^o) dS,$$

where h is the heat transfer coefficient and T^o is the reference temperature. The details of how these quantities are computed are currently omitted in this manual. In general, the quantities are post processed from the fluids temperature field. A surface-based gradient is computed on the boundary face. Nodes on the face augment a heat transfer coefficient field while nodes off the face contribute to a reference temperature.

For radiative heat transfer, the boundary condition applied is as follows:

$$-\kappa \frac{\partial T}{\partial x_j} n_j dS = \epsilon(\sigma T^4 - H) dS,$$

where H is again the irradiation provided by the RTE solve.

If no temperature is specified or an adiabatic line command is used, a zero flux condition is applied.

Species

If a value is specified for each quantity within the wall boundary condition block, a Dirichlet condition is applied. If no values are specified, a zero flux condition is applied.

3.9.3 Atmospheric Boundary Layer Surface Conditions

Monin-Obukhov Theory

Consider atmospheric flow over a flat but non-smooth surface; the coordinate system convention is that flow is along the x -axis, while the z -axis is oriented normal to the surface. The surface layer is the relatively thin layer near the surface where strong wind and temperature gradients exist. Turbulence within this layer can be generated through mechanisms of both shear and thermal convection; the relative contributions of these two mechanisms is determined by the stability state of the atmosphere. The stability state is characterized by the Monin-Obukhov length:

$$L = -\frac{u_\tau^3 \theta_{ref}}{\kappa g (w' \theta')_s};$$

u_τ is the friction velocity, defined as the square root of the magnitude of the Reynolds shear stress at the surface, or

$$u_\tau = \left(\overline{w'u'^2} + \overline{w'v'^2} \right)^{1/4} = \sqrt{\frac{\tau_s}{\rho_s}}$$

θ_{ref} is a reference (virtual potential) temperature associated with the air within the surface layer; for example, the average temperature within the surface layer. $\kappa \approx 0.41$ is the von Karman constant, and g is the acceleration of gravity. $\overline{w' \theta'_s}$ is the surface turbulent temperature flux. Both the turbulent shear stress and turbulent temperature flux are approximately constant within the surface layer.

Applying a gradient diffusion model for the turbulent temperature flux leads to:

$$\overline{w' \theta'_s} = -k_T \frac{\partial \theta}{\partial z}$$

The sign of L is then connected to the sign of the temperature gradient within the surface layer. Three regimes are delineated:

- $\frac{1}{L} > 0$, $\frac{\partial \theta}{\partial z} > 0$, stable stratification
- $\frac{1}{L} = 0$, $\frac{\partial \theta}{\partial z} = 0$, neutral stratification
- $\frac{1}{L} < 0$, $\frac{\partial \theta}{\partial z} < 0$, unstable stratification

Monin-Obukhov theory postulates the following similarity laws for mean velocity parallel to the surface and temperature,

$$\frac{\kappa z}{u_\tau} \frac{\partial \bar{u}_{||}}{\partial z} = \phi_m \left(\frac{z}{L} \right), \quad (3.107)$$

$$\frac{\kappa z u_\tau}{\overline{w' \theta'_s}} \frac{\partial \bar{\theta}}{\partial z} = \phi_h \left(\frac{z}{L} \right), \quad (3.108)$$

where the forms of the non-dimensional functions ϕ_m and ϕ_h are determined from empirical observations. Analytical functions have been fit to the data; these are not given here, rather, we present the integrated form of ((3.107)) and ((3.108)), since these are the forms required by the code implementation.

For neutral stratification, $\phi_m = 1$ and we recover the logarithmic profile for a “fully rough” surface,

$$\bar{u}_{||}(z) = \frac{u_\tau}{\kappa} \ln \frac{z}{z_0}, \quad (3.109)$$

where z_0 is the characteristic roughness height. Note that viscous scaling involving surface viscosity and density properties is not required with this form of the logarithmic profile, since the roughness height is large enough to eliminate the presence of a laminar sublayer and buffer layer.

For stable stratification, the surface layer profiles take the form

$$\bar{u}_{||}(z) = \frac{u_\tau}{\kappa} \left(\ln \frac{z}{z_0} + \gamma_m \frac{z}{L} \right) \quad (3.110)$$

$$\bar{\theta}(z) = \bar{\theta}(z_0) + \frac{\theta_*}{\kappa} \left(\alpha_h \ln \frac{z}{z_0} + \gamma_h \frac{z}{L} \right) \quad (3.111)$$

θ_* is calculated from the temperature flux and friction velocity as $\theta_* = -\frac{\overline{w'\theta'_s}}{u_\tau}$, and γ_m , α_h , and γ_h are constants specified below.

For unstable stratification, the surface layer profiles take the form

$$\bar{u}_{||}(z) = \frac{u_\tau}{\kappa} \left(\ln \frac{z}{z_0} - \psi_m \left(\frac{z}{L} \right) \right) \quad (3.112)$$

$$\bar{\theta}(z) = \bar{\theta}(z_0) + \alpha_h \frac{\theta_*}{\kappa} \left(\ln \frac{z}{z_0} - \psi_h \left(\frac{z}{L} \right) \right) \quad (3.113)$$

where

$$\psi_m \left(\frac{z}{L} \right) = 2 \ln \frac{1+x}{2} + \ln \frac{1+x^2}{2} - 2 \tan^{-1} x + \frac{\pi}{2}, \quad x = \left(1 - \beta_m \frac{z}{L} \right)^{1/4}, \quad (3.114)$$

$$\psi_h \left(\frac{z}{L} \right) = \ln \frac{1+y}{2}, \quad y = \left(1 - \beta_h \frac{z}{L} \right)^{1/2}. \quad (3.115)$$

The constants used in ((3.110)) – ((3.115)) are [Dye74]

$$\kappa = 0.41, \quad \alpha_h = 1, \quad \beta_m = 16, \quad \beta_h = 16, \quad \gamma_m = 5.0, \quad \gamma_h = 5.0.$$

ABL Wall Function

The equations from the preceeding section can be used to formulate a wall function boundary condition for simulation of atmospheric boundary layers. The user-specified inputs to this boundary condition are: roughness length, z_0 , and surface heat flux, $q_s = \rho C_p (\overline{w'\theta'})_s$. The surface layer profile model is evaluated for each surface boundary flux integration point; the wall-normal distance of the “first point off the wall” is taken to be one fourth of the length of the nearest edge intersecting the boundary face. The boundary condition is specified weakly through the imposition of a surface shear stress and surface heat flux.

The procedure for applying the boundary condition is as follows:

1. Determine the stratification state of the boundary layer by calculating the sign of the Monin-Obukhov length scale.
2. Solve the appropriate profile equation, either ((3.109)), ((3.110)), or ((3.112)), for the friction velocity u_τ . For the neutral case, u_τ can be solved for directly. For the stable and unstable cases, u_τ must be solved for iteratively because L appears in these equations and L depends on u_τ .

3. The surface shear stress is calculated as $\tau_s = \rho_s u_\tau^2$. For calculating left-hand-side Jacobian entries, the form (3.116) is used, where ψ' is zero for a neutral profile, $-\gamma_m z/L$ for a stable profile, and $\psi_h(z/L)$ for an unstable profile. The Jacobian entries follow directly from this form.
4. The user specified surface heat flux is applied to the enthalpy equation. Evaluation of surface temperature is not required for the boundary condition specification. However, if surface temperature is required for evaluation of other quantities within the code, the appropriate surface layer temperature profile should be used, either ((3.111)) or ((3.113)).

$$\tau_{si} = \lambda_s u_{||i} = \frac{\kappa \rho u_\tau}{\log(z/z_0) - \psi'(z/L)}, \quad (3.116)$$

3.9.4 Atmospheric Boundary Layer Top Conditions

The abltop option is intended for the upper boundary of atmospheric boundary layer simulations. Currently it has two functions: (1) to provide a fully-automated open boundary condition that allows for inflows and outflows generated by terrain features or obstacles placed within the domain interior, and (2) to allow for mean temperature gradients at the upper boundary. These two features will be described separately below.

Inflow-Outflow Capability

The inflow-outflow capability is facilitated by a potential flow solution on a sub-domain consisting of a thin slab extending from the upper boundary a short distance into the domain. The vertical velocity at the slab bottom face is measured directly from the solution at each time step and this information is used to form a potential flow solution for the entire slab, which effectively provides all three velocity components at the main domain upper boundary. The potential flow solution is achieved at very little computational cost via Fast Fourier Transforms. The mathematical aspects of the boundary condition solution procedure are sketched below.

We start by assuming that the flow near the upper boundary is either laminar or weakly turbulent, as is the norm in atmospheric boundary layer simulations due to the presence of capping inversions (a region of increased stability that effectively terminates the boundary layer). If any turbulence within the slab is weak then the flow is approximately irrotational and any density fluctuations are weak. Furthermore, if the slab is thin (compared with the density scale height ~ 8 km) then the mean density variation is also negligible. Thus we idealize the flow as being irrotational and of constant density, which leads to the following potential flow framework

$$\nabla^2 \phi = 0,$$

$$\vec{u} = \nabla \phi + \langle \vec{u} \rangle,$$

where ϕ is the disturbance potential and $\langle \vec{u} \rangle$ is the mean velocity.

A Fourier transform of Laplace's equation in x and y leads to

$$-(k^2 + l^2)\hat{\phi} + \frac{\partial^2 \hat{\phi}}{\partial z^2} = 0,$$

where k and l are the wavenumber components in the x and y directions respectively, and where $\hat{\phi}$ is the Fourier transform of ϕ .

The bounded solution to the above equation is

$$\hat{\phi}(k, l, z) = \hat{\phi}_0 \exp \left[-\sqrt{k^2 + l^2}(z - z_0) \right],$$

where $\hat{\phi}_0$ is the Fourier transform of ϕ on the horizontal plane $z = z_0$. The vertical velocity component (in Fourier $x - y$ space) is

$$\hat{w}(k, l, z) = \frac{\partial \hat{\phi}}{\partial z} = -\sqrt{k^2 + l^2} \hat{\phi}_0 \exp \left[-\sqrt{k^2 + l^2}(z - z_0) \right].$$

When evaluated at the position $z = z_0$ the above relation yields the following specification for $\hat{\phi}_0$

$$\hat{\phi}_0 = -\frac{1}{\sqrt{k^2 + l^2}} \hat{w}_0.$$

Making use of this result, the solution for $\hat{\phi}$ can be written as

$$\hat{\phi}(k, l, z) = -\frac{1}{\sqrt{k^2 + l^2}} \hat{w}_0 \exp \left[-\sqrt{k^2 + l^2} (z - z_0) \right]$$

If we identify z_0 with the position of the potential flow slab lower face (the sampling plane), then the above prescription for $\hat{\phi}$ yields the solution anywhere in the slab using just knowledge of \hat{w} at its lower boundary. The three velocity components are

$$\begin{aligned} \hat{u}(k, l, z) &= \\ \frac{\partial \hat{\phi}}{\partial x} &= -\frac{ik}{\sqrt{k^2 + l^2}} \hat{w}_0 \exp \left[-\sqrt{k^2 + l^2} (z - z_0) \right] + \langle u \rangle \\ \hat{v}(k, l, z) &= \\ \frac{\partial \hat{\phi}}{\partial y} &= -\frac{il}{\sqrt{k^2 + l^2}} \hat{w}_0 \exp \left[-\sqrt{k^2 + l^2} (z - z_0) \right] + \langle v \rangle \\ \hat{w}(k, l, z) &= \\ \frac{\partial \hat{\phi}}{\partial z} &= \hat{w}_0 \exp \left[-\sqrt{k^2 + l^2} (z - z_0) \right] + \langle w \rangle \end{aligned}$$

Notice that all three velocity components contain a common wavenumber-dependent exponential damping term. This feature indicates that high wavenumber (small-scale) information on the sampling plane is strongly filtered in constructing the velocity field at the upper boundary. Thus the boundary condition velocity field will always be smoother than the velocity on the sampling plane. The filtering also increases with larger separations between the sampling plane and the upper boundary. This fact poses the need to compromise between the additional computational cost of adding a thicker potential flow region and the reduced filtering resulting from a thin slab. We have found that a potential flow region thickness of 10% of the overall computational domain height provides a good compromise between these two competing aspects of the solution.

The solution procedure is rather simple:

1. Measure $w(x, y, z = z_0)$, at the position of the sampling plane.
2. Take the 2D Fourier transform of w to get $\hat{w}_0(k, l)$.
3. Solve for the three velocity components given above using the computational domain upper boundary position for z and the sampling plane position for z_0 .
4. Inverse Fourier transform all three solution components.

While this solution procedure is simple and can be computed at very low cost using Fast Fourier Transforms (FFTs), it does have some obvious restrictions, namely

1. Data must be sampled on a uniform Cartesian grid at a fixed elevation.
2. Periodic or symmetry boundary conditions must exist in the lateral directions.

The simplest way to achieve the first constraint is to add a Cartesian mesh block to the upper portion of the main computational domain. The thickness of the addition can be rather small, say about 10% of the original domain height. The current implementation assumes that a plane at constant elevation with a uniform Cartesian mesh exists near the upper boundary. However, if required in the future, interpolation can be used to sample the solution on a horizontal plane. The `BdyLayerVelocitySampler` class can do this.

While the second constraint may seem restrictive, it allows for any combination of periodic or inflow/outflow boundary conditions in either of the two lateral directions. Half-wave instead of periodic transforms are used

for an inflow/outflow direction. For example, for inflow/outflow in the x direction, we take $u \sim \cos(\pi kx/xL)$, $v \sim \sin(\pi kx/xL)$, $w \sim \sin(\pi kx/xL)$. These prescriptions dictate that $\partial u/\partial x = 0$ at $x = 0$ (the inlet) and at $x = xL$ (the outlet) and that $v = 0$, $w = 0$ at the inlet and outlet. In order to meet these constraints, the terrain should be horizontal for a short distance downstream of the inlet and a short distance upstream of the outlet. The appropriate combination of half-wave and full-wave (periodic) transforms in the two coordinate directions allow for various combinations of inflow/outflow or periodic conditions in these directions.

Temperature Gradient Capability

The temperature almost always varies with height within the atmosphere. The temperature gradient option allows the simulation to be consistent with this fact specifying the rate of change in temperature at the upper boundary.

Implementation Details

The ABL top boundary condition is activated by the keyword `abltop_boundary_condition` in the input file. The following is an example block of an ABL top boundary condition specification:

```
- abltop_boundary_condition: bc_upper
target_name: top
abltop_user_data:
    potential_flow_bc: true
    grid_dimensions: [121, 2, 61]
    horizontal_bcs: [1, -1, 0, 0]
    z_sample: 0.85
    normal_temperature_gradient: 0.01
```

The potential flow and temperature gradient options can be used independently or together. The Boolean input `potential_flow_bc: true` activates the potential flow feature, whereas the presence of the keyword `normal_temperature_gradient: value` activates the temperature gradient feature. If only the temperature gradient feature is called for, then symmetry conditions (i.e. a slip wall) are used continuity and momentum boundary conditions at the upper boundary. If only the potential flow boundary condition is called for then the condition $d\langle T \rangle/dz = 0$ is used at the upper boundary (where $\langle \rangle$ indicates a horizontal average).

If the potential flow option is selected, the user must then specify three additional user inputs: `grid_dimensions`, `horizontal_bcs` and `z_sample`. The `grid_dimensions` input specifies the number of mesh points in the three coordinate directions for the structured Cartesian portion of the domain containing the potential flow slab. This may be the entire domain, or just a sub-portion of it. The current implementation assumes that the grid points within the structured Cartesian region are tagged with a mesh index flag that indicates the relative position of each point within the structured mesh system. For example, if $(i_{max}, j_{max}, k_{max})$ are the grid dimensions in the three coordinate directions, then the mesh index for the point (i, j, k) is $k*i_{max}*j_{max} + j*i_{max} + i$. The `abl_mesh` program will include the mesh index tag in the `exodus` grid file.

The `horizontal_bcs` input specifies the lateral boundary conditions in use in the i and j grid directions. Inflow is specified with +1, outflow -1, and periodic with zero. Thus in the example above, inflow is used at the $i=0$ boundary, outflow at the $i=i_{max}$ boundary, and periodic conditions are used in the j direction.

The input `z_sample` specifies the elevation of the sampling plane. If this input is missing, the default position of 90% of the distance between the lower and upper boundary will be used.

3.9.5 Turbulent Kinetic Energy, k_{sgs} LES model

When the boundary layer is assumed to be resolved, the natural boundary condition is a Dirichlet value of zero, $k_{sgs} = 0$.

When the wall model is used, a standard wall function approach is used with the assumption of equal production and dissipation.

The turbulent kinetic energy production term is consistent with the law of the wall formulation and can be expressed as,

$$P_{kw} = \tau_w \frac{\partial u_{\parallel}}{\partial y}. \quad (3.117)$$

The parallel velocity, u_{\parallel} , can be related to the wall shear stress by,

$$\tau_w \frac{u^+}{y^+} = \mu \frac{u_{\parallel}}{Y_p}. \quad (3.118)$$

Taking the derivative of both sides of Equation (3.118), and substituting this relationship into Equation (3.117) yields,

$$P_{kw} = \frac{\tau_w^2}{\mu} \frac{\partial u^+}{\partial y^+}. \quad (3.119)$$

Applying the derivative of the law of the wall formulation, Equation (3.94), provides the functional form of $\partial u^+ / \partial y^+$,

$$\frac{\partial u^+}{\partial y^+} = \frac{\partial}{\partial y^+} \left[\frac{1}{\kappa} \ln(Ey^+) \right] = \frac{1}{\kappa y^+}. \quad (3.120)$$

Substituting Equation (3.94) within Equation (3.119) yields a commonly used form of the near wall production term,

$$P_{kw} = \frac{\tau_w^2}{\rho \kappa u_{\tau} Y_p}. \quad (3.121)$$

Assuming local equilibrium, $P_k = \rho \epsilon$, and using Equation (3.121) and Equation (3.95) provides the form of wall shear stress is given by,

$$\tau_w = \rho C_{\mu}^{1/2} k. \quad (3.122)$$

Under the above assumptions, the near wall value for turbulent kinetic energy, in the absence of convection, diffusion, or accumulation is given by,

$$k = \frac{u_{\tau}^2}{C_{\mu}^{1/2}}. \quad (3.123)$$

This expression for turbulent kinetic energy is evaluated at the boundary faces of the exposed wall boundaries and is area-assembled to the nodal value for use in a Dirichlet condition.

Turbulent Kinetic Energy and Specific Dissipation SST Low Reynolds Number Boundary conditions

For the turbulent kinetic energy equation, the wall boundary conditions follow that described for the k_{sgs} model, i.e., $k = 0$.

A Dirichlet condition is also used on ω . For this boundary condition, the ω equation depends only on the near-wall grid spacing. The boundary condition is given by,

$$\omega = \frac{6\nu}{\beta_1 y^2},$$

which is valid for $y^+ < 3$.

Turbulent Kinetic Energy and Specific Dissipation SST High Reynolds Number Boundary conditions

The high Reynolds approach uses the law of the wall assumption and also follows the description provided in the wall modeling section with only a slight modification in constant syntax,

$$k = \frac{u_\tau^2}{\sqrt{\beta^*}}. \quad (3.124)$$

In the case of ω , an analytic expression is known in the log layer:

$$\omega = \frac{u_\tau}{\sqrt{\beta^*} \kappa y},$$

which is independent of k . Because all these expressions require y to be in the log layer, they should absolutely not be used unless it can be guaranteed that $y^+ > 10$, and $y^+ > 25$ is preferable. Automatic blending is not currently supported.

Solid Stress

The boundary conditions applied are either force provided by a static pressure,

$$F_i^n = \int \bar{P} n_i dS, \quad (3.125)$$

or a Dirichlet condition, i.e., $u_i = u_i^{spec}$, on the displacement field. Above, F_i^n is the force for component i due to a prescribed [static] pressure.

Intensity

The boundary condition for each intensity assumes a grey, diffuse surface as,

$$I(s) = \frac{1}{\pi} [\tau \sigma T_\infty^4 + \epsilon \sigma T_w^4 + (1 - \epsilon - \tau) K]. \quad (3.126)$$

3.9.6 SST RANS Model for Atmospheric Boundary Layer

The following boundary conditions simulate the Atmospheric Boundary Layer, as described in Bautista, [Bau11] and [BDM15]. The Nalu-Wind SST RANS implementation matches the Monin-Obukhov profile when used with the model constants from Table-A I-1 (Boundreault, 2011) in [Bau11] and the meshing method described in [BDM15]. The mesh described in [BDM15] gives the Monin-Obukhov profile for roughness height 0.1. When the roughness height is decreased, the mesh must be refined near the wall. For example, for the [BDM15] ABL test case using roughness height 0.001 instead of 0.1, the mesh size needs to be halved near the wall.

The k and ω boundary conditions are the same as in the *k-omega SST boundary conditions*:

$$k = \frac{u_\tau^2}{\sqrt{\beta^*}},$$

and

$$\omega = \frac{u_\tau}{\sqrt{\beta^*} \kappa y}.$$

The momentum boundary condition is a no-slip Dirichlet condition, $u_i = 0$, as described in the *momentum wall boundary conditions*.

The streamwise and spanwise boundary conditions are periodic, as described in the *periodic boundary conditions*.

The k , ω , and u wall boundary conditions are set in the input file by specifying a wall boundary condition with `RANS_abl_bc`. The input file must also specify a height and the velocity at that height with `reference_height` and `reference_velocity`.

```
realms:
- name: fluidRealm
  boundary_conditions:
- wall_boundary_condition: bc_lower
  wall_user_data:
    RANS_abl_bc: yes
    reference_velocity: 6.6
    reference_height: 90.0
```

This height, h , and velocity, u_h , could, for example, be the hub height of a wind turbine and the velocity measured at that height.

The input file should also include a momentum source term, `momentum`.

```
realms:
- name: fluidRealm
  solution_options:
    options:
- source_terms:
    momentum:
- body_force
- source_term_parameters:
    momentum: [0.0003547, 0.0, 0.0]
```

The momentum source term, dp/dx , is calculated by balancing this pressure gradient with the wall shear stress, τ_w .

$$\frac{dp}{dx}V = \tau_w A$$

where V is the volume of the domain and A is the area of the domain that touches the ground. Cancelling the length and width of the domain and dividing by the height of the domain, H , gives

$$\frac{dp}{dx} = \frac{\rho u_\tau^2}{H}$$

u_τ is calculated from the Monin-Obukhov profile,

$$u_\tau = \frac{u_h \kappa}{\log((h + z_0)/z_0)}$$

where κ is the Von K{ 'a}rm{ 'a}n constant and z_0 is the roughness height.

3.9.7 Open Boundary Condition

Open boundary conditions require far more care. In general, open bcs are assembled by iterating faces and the boundary integration points on the exposed face. The parent element is also required since oftentimes gradients are used (for momentum). For an open boundary condition the flow can either leave or enter the domain depending on what the computed mass flow rate at the exposed boundary integration point is. Two options are available computing the velocity of the entrained flow—either the normal velocity at the integration point is used or a specified normal velocity is used. The tangential components are always specified.

Continuity

For continuity, the boundary mass flow rate must also be computed. This value is stored and used for the other equations that require advection. The same formula is used for the pressure-stabilized mass flow rate. However, the

local pressure gradient for each boundary contribution is based on the difference between the interior integration point and the user-specified pressure which takes on the boundary value. This can optionally be modified to be a “total pressure”—removing the kinetic energy associated with entrainment at the open. The interior integration point is determined by linear interpolation. For CVFEM, full elemental averaging is used while in EBVC discretization, the midpoint value between the nearest node and opposing node to the boundary integration point is used. In both discretization approaches, non-orthogonal corrections are required. This procedure has been very important for stability for CVFEM tet-based meshes where a natural non-orthogonality exists between the boundary and interior integration point.

For wind energy applications, the usage of the standard open boundary mass flow rate expression, which includes pressure contributions, is not appropriate due to complex temperature/buoyancy specifications. In these cases, a global correction algorithm is supported. Specifically, pressure terms are dropped at the open boundary mass flow rate expression in favor of a pre-processing algorithm that uniformly distributes the continuity mass flow rate (and possible density accumulation) “error” over the entire set of open boundary conditions. The global correction scheme may perform well with single open boundary condition specification, e.g., multiple inflows with a single open location, however, it is to be avoided if the flow leaving the domain is complex in that a simulation includes multiple open boundary conditions. A complex situation might be an open jet with entrainment from the side (open boundary that allows for inflow) and a top open that allows for outflow. However, a routine case might be a backward facing step with a single inflow, side periodic, top wall and open boundary. Not that the ability for the continuity solve to be well conditioned may require an interior Dirichlet on pressure as the open pressure specification for the global correction algorithm is lacking. In most cases, a Dirichlet condition is not actually required as the nullspace of the continuity system may not be found in the solve.

Momentum

For momentum, the normal component of the stress is subtracted out we subtract out the normal component of the stress. The normal stress component for component i can be written as $F_k n_k n_i$. The tangential component for component i is simply, $F_i - F_k n_k n_i$. As an example, the tangential viscous stress for component x is,

$$F_x^T = F_x - (F_x n_x + F_y n_y) n_x,$$

which can be written in general component form as,

$$F_i^T = F_i(1 - n_i n_i) - \sum_{j \neq i} F_j n_i n_j.$$

Finally, the normal stress contribution is applied based on the user specified pressure,

$$F_i^N = P^{Spec} A_i.$$

For CVFEM, the face gradient operators are used for the thermal stress terms. For EBVC discretization, from the boundary integration point, the nearest node (the “Right” state) is used as well as the opposing node (the “Left” state). The nearest node and opposing node are used to compute gradients required for any derivatives. This equation follows the standard gradient description in the diffusion section with non-orthogonal corrections used. In this formulation, the area vector is taken to be the exposed area vector. Non-orthogonal terms are noted when the area vector and edge vector are not aligned.

For advection, if the flow is leaving the domain, we simply advect the nearest nodal value to the boundary integration point. If the flow is coming into the domain, we simply confine the flow to be normal to the open boundary integration point area vector. The value entrained can be the nearest node or an upstream velocity value defined by the edge midpoint value or by a specified value.

Mixture Fraction, Enthalpy, Species, k_{sgs} , k and ω

Open boundary conditions assume a zero normal gradient. When flow is entering the domain, the far-field user supplied value is used. Far field values are used for property evaluations. When flow is leaving the domain, the flow is advected out consistent with the choice of interior advection operator.

3.9.8 Symmetry Boundary Condition

Continuity, Mixture Fraction, Enthalpy, Species, k_{sgs} , \mathbf{k} and ω

Zero diffusion is applied at the symmetry bc.

Momentum

A symmetry boundary is one that is described by removal of the tangential stress. Therefore, only the normal component of the stress is applied:

$$F_x^n = (F_x n_x + F_y n_y) n_x,$$

which can be written in general component form as,

$$F_i^n = F_j n_j n_i.$$

Specified Boundary-Normal Temperature Gradient Option

The standard symmetry boundary condition applies zero diffusion at the boundary for scalar quantities, which effectively results in those scalars having a zero boundary-normal gradient. There are situations, especially for atmospheric flows in which the user may desire a finite boundary-normal gradient of temperature. For example, the atmospheric boundary layer is often simulated with a stably stratified capping inversion in which the temperature linearly increases with height all the way to the upper domain boundary. We apply symmetry conditions to this upper boundary for momentum, but we specify the boundary-normal temperature gradient on this boundary to match the capping inversion's gradient.

This is an option in the symmetry boundary condition specification, which appears in the input file as:

```
- symmetry_boundary_condition: bc_upper
  target_name: upper
  symmetry_user_data:
    normal_temperature_gradient: -0.003
```

In this example, the temperature gradient normal to the symmetry boundary is set to -0.003 K/m, where the boundary-normal direction is pointed into the domain.

Nalu-Wind does not solve a transport equation for temperature directly, but rather it solves one for enthalpy. Therefore, the boundary-normal temperature gradient condition is applied internally in the code through application of a compatible heat flux,

$$q_n = -\kappa_{eff} c_p \frac{\partial T}{\partial n}$$

where q_n is the heat flux at the boundary, κ_{eff} is the effective thermal diffusivity (the molecular and turbulent parts), c_p is the specific heat, and $\partial T / \partial n$ is the boundary-normal temperature gradient.

3.9.9 Periodic Boundary Condition

A parallel multiple-periodic boundary condition is supported. Mappings are created between master/slave surface node pairs. The node pairs are obtained from a parallel search and are expected to be unique. The node pairs are used to map the slave global id to that of the master. This allows the linear system to include matrix rows for only a subset of the overall set of nodes. Moreover, a periodic assembly for assembled quantities is managed via: $m+ = s$ and $s = m$, where m and s are master/slave nodes, respectively. For each parallel assembled quantity, e.g., dual volume,

turbulence quantities, etc., this procedure is used. Periodic boxes and periodic couette and channel flow have been simulated in this code base. Two forms of parallel searches exist and are supported (one through the Boost TPL and another through the STK Search module).

3.9.10 Non-conformal Boundary Condition

A surface-based approach based on a DG method has been discussed in the 2010 CTR summer proceedings by Domino, [Dom10]. Both the edge- and element-based formulation currently exists in the code base using the CVFEM and EBVC approaches.




Fig. 3.6: Two-block example with one common surface, Γ_{AB} .

Consider two domains, A and B , which have a common interface, Γ_{AB} , and a set of interfaces not in common, $\Gamma \setminus \Gamma_{AB}$ (see Figure Fig. 3.6), and assume that the solution of the time-dependent advection/diffusion equation is to be solved in both domains. Each domain has a set of outwardly pointing normals. In this cartoon, the interface is well resolved, although in practice this may not be the case.

An interior penalty approach is constructed at each integration point at the exposed surface set. The numerical flux for a general scalar ϕ is constructed at the current integration point which is based on the current (A) and opposing (B) elemental contributions,

$$\int \hat{Q}^A dS = \int \left[\frac{(q_j^A n_j^A - q_j^B n_j^B)}{2} + \lambda^A (\phi^A - \phi^B) \right] dS^A + \dot{m}^A \frac{(\phi^A + \phi^B)}{2} + \eta \frac{|\dot{m}^A|}{2} (\phi^A - \phi^B), \quad (3.127)$$

where q_j^A and q_j^B are the diffusive fluxes computed using the current and opposing elements and normals are outward facing. The penalty coefficient λ^A contains the diffusive contributions averaged over the two elements,

$$\lambda^A = \frac{(\Gamma^A/L^A + \Gamma^B/L^B)}{2}. \quad (3.128)$$

Above, Γ^k is the diffusive flux coefficient evaluated at current and opposing element location, respectively, and L^k is an elemental length scale normal to the surface (again for current and opposing locations, A and B). When upwinding is activated, the value of η is unity.

As written in Equation (3.127), the default convection and diffusion term is a Galerkin approach, i.e., equally averaged between the current and opposing face. The standard advection term is given by,

$$\int \rho \hat{u}_j \phi n_j dS. \quad (3.129)$$

For surface A, the form is as follows:

$$\int \rho \hat{u}_j^A \phi n_j^A dS^A = \dot{m}^A \frac{\phi^A + \phi^B}{2}, \quad (3.130)$$

with the nonconformal mass flow rate given by,

$$\dot{m}^A = \left[\frac{(\rho u_j^A + \gamma(\tau G_j^A p - \tau \frac{\partial p^A}{\partial x_j})) n_j^A - (\rho u_j^B + \gamma(\tau G_j^B p - \tau \frac{\partial p^B}{\partial x_j})) n_j^B}{2} + \lambda^A (p^A - p^B) \right] dS^A. \quad (3.131)$$

In the above set of expressions, the consistent definition of \hat{u}_j , i.e., the convecting velocity including possible pressure stabilization terms, is retained.

As with the interior advection scheme, the mass flow rate involves pressure stabilization terms. The value of γ defines whether or not the full pressure stabilization terms are included in the mass flow rate expression. Equation (3.131) also forms the continuity nonconformal boundary contribution.

With the substitution of η to be unity, the effective convective term is as follows:

$$\int \rho \hat{u}_j \phi n_j^A dS^A = \frac{(\dot{m}^A + |\dot{m}^A|) \phi^A + (\dot{m}^A - |\dot{m}^A|) \phi^B}{2}. \quad (3.132)$$

Note that this form reduces to a standard upwind operator.

Since this algorithm is a dual pass approach, a numerical flux can be written for the integration point on block B,

$$\int \hat{Q}^B dS = \int \left[\frac{(q_j^B n_j^B - q_j^A n_j^A)}{2} + \lambda^B (\phi^B - \phi^A) \right] dS^A + \dot{m}^B \frac{(\phi^B + \phi^A)}{2} + \eta \frac{|\dot{m}^B|}{2} (\phi^B - \phi^A). \quad (3.133)$$

As with Equation (3.133), \dot{m}^B (see Equation (3.134)) is of similar form to \dot{m}^A ,

$$\dot{m}^B = \left[\frac{(\rho u_j^B + \gamma(\tau G_j^B p - \tau \frac{\partial p^B}{\partial x_j})) n_j^B - (\rho u_j^A + \gamma(\tau G_j^A p - \tau \frac{\partial p^A}{\partial x_j})) n_j^A}{2} + \lambda^A (p^B - p^A) \right] dS^B. \quad (3.134)$$

For low-order meshes with curved surface, faceting will occur. In this case, the outward facing normals may not be (sign)-unity factors of each other. In this case, it may be advantageous to define the opposing outward normal as, $n_j^B = -n_j^A$.

Domino, [Dom10] provided an overview of a FEM fluids implementation. In such a formulation, the interior penalty term appears, i.e.,

$$\int_{\Gamma_{AB}} \frac{\partial w^A}{\partial x_j} n_j \lambda (\phi^A - \phi^B) d\Gamma,$$

and

$$\int_{\Gamma_{BA}} \frac{\partial w^B}{\partial x_j} n_j \lambda (\phi^B - \phi^A) d\Gamma.$$

Although the sign of this term is often debated in the literature, the above set of expressions acts to increase penalty term stencil to include the full element contribution. As the CVFEM uses a piecewise-constant test function, this term is currently neglected.

Average fluxes are computed based on the current and opposing integration point locations. The appropriate DG terms are assembled as boundary conditions first with block A integration points as *current* (integrations points for block B are *opposing*) and then with block B integration points as *current* (surfaces for block A are, therefore, *opposing*). Figure Fig. 3.6 graphically demonstrates the procedure in which integration point values of the flux and penalty term are computed on the block A surface and at the projected location of block B.

A parallel search is conducted to project the current integration point location to the opposing element exposed face. The search, therefore, provides the isoparametric coordinates on the opposing element. Elemental shape functions and

shape function derivatives are used to construct the numerical flux for both the edge- and element-based scheme. The location of the Gauss points on the current element are either the Gauss Labatto or Gauss Legendre locations (input file specification). For each equation (momentum, continuity, enthalpy, etc.) the numerical flux is computed at each exposed non-conformal surface.

As noted, for most equations other than continuity and heat condition, the numerical flux includes advection and diffusion contributions. The diffusive contribution is easily provided using elemental shape function derivatives at the current and opposing surface.




Fig. 3.7: Description of the numerical flux calculation for the DG algorithm. The value of fluxes and penalty values on the current block (A) and the opposing block (B) are used for the calculation of numerical fluxes. $\tilde{\varphi}$ represents the projected value.

Above, special care is taken for the value of the mass flow rate at the non-conformal interface. Also, note that the above written form does not upwind the advective flux, although the code allows for an upwinded approach. In general, the advective term contains contributions from both elements identified at the interface, specifically.

The penalty coefficient for the mass flow rate at the non-conformal boundary points is again a function of the blended inverse length scale at the current and opposing element surface location. The form of the mass flow rate above provides the continuity contribution and the form of the mass flow rate used in the scalar non-conformal flux contribution.

The full connectivity for element integration and opposing elements is within the linear system. As such, for sliding mesh configurations, the linear system connectivity graph changes each time step. Recent prototyping of the dG-based and the overset scheme has allowed this method to be used across both disparate low-order topologies (see Figure Fig. 3.8 and Figure Fig. 3.9).

3.10 Overset

Nalu-Wind supports simulations using an overset mesh methodology to model complex geometries. Currently the codebase supports two approaches to determine overset mesh connectivity:

1. Overset mesh hole-cutting algorithm based on native STK search routines, and
2. Hole-cutting and donor/reception determination using the **TIOGA** (Topology Independent Overset Grid Assembly) TPL.

The native STK based overset grid assembly (OGA) requires no additional packages, but is limited to simple geometries where the search and hole-cutting procedure works only simple rectangular boundaries (for the inner mesh) that are aligned along the major axes. On the other hand, TIOGA based hole cutting is capable of performing overset grid assembly on arbitrary mesh geometries and orientation, supports generalized mesh motion, and can determine




Fig. 3.8: A low-order and high-order block interface ($P=1$ quad4 and $P=2$ quad9) for a MMS temperature solution. In this image, the inset image is a close-up of the nodal Ids near the interface that highlights the quad4 and quad9 interface.




Fig. 3.9: Discontinuous Galerkin non-conformal interface mixed topology (hex8/tet4).

donor/recipient status with multiple meshes overlapping in the same space. A specific use-case for the need to perform OGA on multiple meshes is the simulation of a wind turbine in an atmospheric boundary layer, where the turbine blade, nacelle, and the background ABL mesh might all overlap near the rotor hub.

3.10.1 Overset Grid Assembly using Native STK Search

The overset descriptions begins with the basic background mesh (block 1) and overset mesh (block 2) depicted in Figure Fig. 3.10. Also shown in this figure is the reduction outer surface of block 2 (light blue). Elements within this reduced overset block will be determined by a parallel search. The collection of elements within this bounding box will be skinned to form a surface on which orphan nodes are placed. Elements within this volume are set in a new internally managed inactive block. These mesh entities are fully removed from the overall matrix for each dof. Elements within this volume are provided a masking integer element varibale of unity to select out of the visualization tool. Therefore, orphan nodes live at the external boundary of block 2 and along the reduced surface. The parallel search provides the mapping of orphan node and owning element from which the state can be constructed.




Fig. 3.10: Two-block use case describing background mesh (block 1) and overset mesh (block 2).

After the full search and overset initialization, this simple example yields the original block 1 and 2, the newly created inactive block 3, the original surface of the overset mesh and the new skinned surface (101) of the inactive block (Figure Fig. 3.11).

A simple heat conduction example is provided in Figure Fig. 3.12 where the circular boundary is set at a temperature of 500 with all external boundaries set to adiabatic.

As noted before, every orphan node lies within an owning element. Sufficient overlap is required to make the system well posed. A fully implicit procedure is provided by writing the orphan node value as a linear constraint of the owning element (Figure Fig. 3.13).

For completeness, the constraint equation for any dof ϕ^o is simply,

$$\phi^o - \sum N_k \phi_k = 0. \quad (3.135)$$

As noted, full sensitivities are provided in the linear system by constructing a row entry with the columns of the nodes within the owning element and the orphan node itself.

Finally, a mixed hex/tet mesh configuration example (overset mesh is tet while background is hex) is provided in Figure Fig. 3.14.




Fig. 3.11: Three-block and two surface, post over set initialization.




Fig. 3.12: A simple heat conduction example providing the overset mesh and donor orphan nodes.




Fig. 3.13: Orphan nodes for background and overset mesh for which a fully implicit constraint equation is written.




Fig. 3.14: Flow past a three-dimensional sphere using a hybrid topology (hex/tet) mesh configuration.

3.10.2 Overset Grid Assembly using TIOGA

Topology Independent Overset Grid Assembler (TIOGA) is an open-source connectivity package that was developed as an academic/research counterpart for PUNDIT (the overset grid assembler used in HPCMP CREATE™ A/V program and HELIOS). The base library has been modified to remove the limitation where each MPI rank could only own one mesh block. The code has been extended to handle multiple mesh blocks per MPI rank to support Nalu-Wind's mesh decomposition strategies.

TIOGA uses a different nomenclature for overset mesh assembly. A brief description is provided here to familiarize users with the differences in nomenclature used in the previous section. When determining overset connectivity, TIOGA ends up assigning `IBLANK` values to the nodes in a mesh. The `IBLANK` field is an integer field that determines the status of the node which can be one of three states:

field point

A field point is a node that behaves as a normal mesh point, i.e., the equations are solved on these nodes and the linear system assembly proceeds as normal. The *field points* are indicated by an `IBLANK` value of 1.

fringe point

A fringe point is a receptor on the receiving mesh where the solution field is mapped from the donor element. A fringe point is indicated by an `IBLANK` value of -1. Fringe points are how information is transferred between the participating meshes. Note that fringe points are referred to as *orphan points* in the STK based overset description.

hole point

A hole point is a node on a mesh that occurs inside a solid body being modeled in another mesh. These points have no valid solution for the equations solved and should not participate in the linear system.

In addition to the `IBLANK` status, the following terms are useful when using TIOGA

donor element

The element that is used to *interpolate* field data from donor mesh to a recipient mesh. While TIOGA provides flow interpolation routines, the current implementation in Nalu-Wind uses the `MasterElement` classes in Nalu-Wind to maintain consistency between the STK and the TIOGA overset implementations.

orphan points

The term orphan point is used differently in TIOGA than the STK based overset implementation. TIOGA refers to nodes as orphan points when there it cannot find a suitable donor element for those nodes that are considered fringe points. This can happen when the nodes on the enclosing element are themselves labeled fringe points.

Unlike the STK based hole cutting approach, that uses predefined bounding boxes to determine donor/receptor locations, TIOGA uses the element volume as the metric to determine the field and fringe points. The high level hole cutting algorithm can be described in the following steps:

- Determine and tag hole points that are fully enclosed within solid bodies, tag neighboring points to be fringe points.
- Determine and flag all mandatory fringe points, e.g., embedded boundaries of interior meshes.
- Determine fringe locations for the exterior meshes where information is transferred back from interior meshes to the exterior/background mesh.

In the current integration, only the hole-cutting and donor/receptor information is processed by the TIOGA library. The linear system assembly, specifically the constraint equations for the fringe points are managed by the same classes that are used with the native STK hole-cutting approach.

Figure Fig. 3.15 shows the field and fringe points as determined by TIOGA during the hole-cutting process. The central white region shows the mesh points of the interior mesh. The salmon colored region shows the overlapping field points where the flow equations are solved on both participating meshes. The green-ish boundary shows the mandatory fringe points for the interior mesh along its outer boundary. The interior boundary of the overlap region are the fringe points for the background mesh where information is transferred from the interior mesh. The extent of the overlap region is determined by the number of element layers necessary to ensure adequate separation between the fringe boundaries on the participating meshes.




Fig. 3.15: TIOGA overset hole cutting for a rotated internal mesh configuration showing the field and fringe locations.

Figure `tioga-overset-cyl` shows the resulting overset assembly for cylinder mesh and a background mesh with an intermediate refinement zone. The hole points (inside the cylinder) have been removed from the linear system for both the intermediate and background mesh. The magenta region shows the overlap of field points of the cylinder and the intermediate mesh. And the yellow region shows the overlap between the background and the intermediate mesh.

Figures Fig. 3.17 and Fig. 3.18 shown the velocity and vorticity contours for the flow past a cylinder simulated using the overset mesh methodology with TIOGA overset connectivity.

3.11 Property Evaluations

Property specification is provided in the material model section of the input file. Unity Lewis number assumptions for diffusive flux coefficients for mass fraction and enthalpy are assumed.

3.11.1 Density

At present, property evaluation for density is given by constant, single mixture fraction-based, HDF5 tables, or ideal gas. For ideal gas, we support a non-isothermal, non-uniform and even an acoustically compressible form.




Fig. 3.16: Overset mesh configuration for simulating flow past a cylinder using a three mesh setup: near-body, body-fitted cylinder mesh, intermediate refined mesh, and coarse background mesh.




Fig. 3.17: Velocity field for a flow past cylinder simulating using an overset mesh methodology with TIOGA mesh connectivity approach.




Fig. 3.18: Vorticity field for a flow past cylinder simulating using an overset mesh methodology with TIOGA mesh connectivity approach.

3.11.2 Viscosity

Property evaluation for viscosity is given by constant, single mixture fraction-based, simple tables or Sutherland's three coefficient as a function of temperature. When mixtures are used, either by reference or species transport, only a mass fraction-weighted approach is used.

3.11.3 Specific Heat

Property evaluation for specific heat is either constant or two-band standard NASA polynomials; again species composition weighting are used (either transported or reference).

3.11.4 Lame Properties

Lame constants are either of type constant or for use in mesh motion/smoothing geometric whereby the values are inversely proportional to the dual volume.

3.12 Coupling Approach

The classic low Mach implementation uses an incremental approximate pressure projection scheme in which nonlinear convergence is obtained using outer Picard loops. Recently a full study on coupling approaches has been conducted using ASC Algorithm funds. In this project, coupling methods ranging from fully implicit, fully coupled equal order pressure/velocity interpolation with pressure stabilization to explicit advection/diffusion pressure projection schemes. A brief summary of the results follows.

Specifically, five algorithms were considered and are as follows:

1. A monolithic scheme in which advection and diffusion are implicit using full analytical sensitivities,

2. Monolithic momentum solve with implicit advection/diffusion in the context of a fourth order stabilized incremental pressure projection scheme,
3. Monolithic momentum solve with explicit advection; implicit diffusion in the context of a fourth order stabilized incremental pressure projection scheme,
4. Segregated momentum solve with implicit advection/diffusion in the context of a fourth order stabilized incremental pressure projection scheme, and
5. Explicit momentum advection/diffusion predictor/corrector scheme in the context of a second order stabilized pressure-free approximate projection scheme.

Each of the above algorithms has been run in the context of a transient uniform flow low Mach flow. The emphasis of this project is transient flows. As such, the numbers below are to be cast in this context. If steady flows are desired, conclusions may be different. The slowdown of each implementation is relative to the core low Mach algorithm, i.e., algorithm (4) above. Numbers less than unity represent a speed-up whereas numbers greater than unity represent a slow down: 1) 3.4x, 2) 1.2x, 3) 0.6x, 4) 1.0x, 5) 0.7x.

The above runs were made using a time step that corresponded to a CFL of slightly less than unity. In this particular flow, a transitionally turbulent open jet, the diffusion time scale stability limit was not a factor. In other words, there existed no detailed boundary layer at the wall bounded flow at the ground plane. Results for a Reynolds number of 45000 back step also are similar to the above jet results.

In general, although a mixture of implicit diffusion and explicit advection seem to be the winning combination, this scheme is very sensitive to time step and must be used by an educated user. In general, the conclusions are, thus far, that the standard segregated pressure projection scheme is preferred.

The algorithm implemented in Nalu-Wind is a fourth order approximate projection scheme with monolithic momentum coupling. Evaluation of a predictor/corrector approach for reacting flow is anticipated in the late FY15 time frame.

3.12.1 Errors due to Splitting and Stabilization

As noted in many of our papers, the error in the above method can be written in block form (let's relax the variable density nuance - or simply fold these extra terms into our operators). Here we specifically partition error into both splitting (the pressure projection aspect of the alg that factorizes the fully coupled system) and pressure stabilization. Note that when we run fully coupled simulation with the same pressure stabilization algorithm, the answers converge to the same result.

Below, also forgive the specific definitions of τ . In general, they represent a choice of projection and stabilization time scales. Finally, the Laplace operator, e.g., \mathbf{L}_2 , have the τ 's built into them.

$$\begin{bmatrix} \mathbf{A} & \mathbf{G} \\ \mathbf{D} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{u}^{n+1} \\ p^{n+1} \end{bmatrix} = \begin{bmatrix} \mathbf{f} \\ 0 \end{bmatrix} + \begin{bmatrix} (\mathbf{I} - \tau \mathbf{A}) \mathbf{G} (p^{n+1} - p^n) \\ \epsilon(\mathbf{L}_1, \tau_1, \mathbf{D}, \mathbf{G}) \end{bmatrix} \quad (3.136)$$

where the error term that appears for the discrete continuity solve is given by,

$$\begin{aligned} \epsilon(\mathbf{L}_1, \tau_1, \mathbf{D}, \mathbf{G}) = & ((\mathbf{L}_1 - \mathbf{D} \tau_1 \mathbf{G}) \\ & - (\mathbf{L}_2 - \mathbf{D} \tau_2 \mathbf{G}))(p^{n+1} - p^n) \\ & + (\mathbf{L}_2 - \mathbf{D} \tau_2 \mathbf{G}) p^{n+1} \end{aligned} \quad (3.137)$$

For the sake of this write-up, let $\mathbf{L}_1 = \mathbf{L}_2$ and $\tau_2 = \tau_1$.

3.13 Time discretization

Time integrators range from simple backward Euler or a second order three state scheme, BDF2.

A general time discretization approach can be written as,

$$\int \frac{\partial \rho \phi}{\partial t} dV = \int \frac{(\gamma_1 \rho^{n+1} \phi^{n+1} + \gamma_2 \rho^n \phi^n + \gamma_3 \rho^{n-1} \phi^{n-1})}{\Delta t} dV$$

where γ_i represent the appropriate factors for either Backward Euler or a three-point BDF2 scheme. In both discretization approaches, the value for density and other dofs are evaluated at the node. As such, the time contribution is a lumped mass scheme with the volume simply the dual volume. The topology over one loops to assemble system is simply the node. Although CVFEM affords the use of a consistent mass matrix, this scheme is not used at present.

3.14 Multi-Physics

The equation set required to support the energy sector is already represented as a multiphysics application. However, in some common cases of coupling including conjugate heat transfer and coupling to participating media radiation, an operator split method may be preferred. The general concept is to define multiple Nalu-Wind Realms that each own the mesh on which the particular physics is solved. Surface- and volume-based couplings are supported through linear interpolation of the coupling parameters.

A typical CHT application involves the coupling of a thermal response and fluid transport. The coupling occurs between the surface that shares the thermal equation and static enthalpy equation. Moreover, coupling to a PMR solve is a volume-based coupling. Multiple Realms are supported with multiple transfers.

In Nalu-Wind, the method to achieve coupling in CHT or RTE coupled systems is through the usage of the STK Transfer module. This allows for linear interpolation between disparate meshes. Advanced conservative transfers are being evaluated, however, are not yet implemented in the code base. In general, the STK Transfer interface allows for this design point.

For FSI, the usage of the transfer module is also expected.

3.15 Wind Energy Modeling

Wind energy analysis is the primary application area for the Nalu-Wind development team. This section describes the theoretical basis of Nalu-Wind from a wind energy perspective, using nomenclature familiar to wind energy experts and mapping it to Nalu-Wind concepts and nomenclature described in previous sections. Hopefully, this will provide an easier transition for users familiar with WRF and SOWFA to Nalu-Wind.

In order to evaluate the energy output and the structural loading on wind turbines, the code must model: 1. the incoming turbulent wind field across the entire wind farm, and 2. the evolution of turbine wakes in turbulent inflow conditions and their interaction with the downstream turbines. First, the governing equations with all the terms necessary to model a wind farm are presented with links to implementation and verification details elsewhere in the theory and/or verification manuals. A brief description of Nalu-Wind's numerical discretization schemes is presented next. This is followed by a brief discussion of the boundary conditions used to model atmospheric boundary layer (ABL) flows with or without wind turbines (currently modeled as actuator sources within the flow domain).

Currently Nalu-Wind supports two types of wind simulations:

Precursor simulations

Precursor simulations are used in wind applications to generate time histories of turbulent ABL inflow profiles that are used as inlet conditions in subsequent wind farm simulations. The primary purpose of these simulations are to trigger turbulence generation and obtain velocity and temperature profiles that have *converged* to a statistic equilibrium.

Wind farm simulation with turbines as actuator sources

In this case, the wind turbine blades and tower are modeled as actuator source terms by coupling to the [OpenFAST](#) libraries. Velocity fields are sampled at the blade and tower control points within the Nalu-Wind domain and the blade positions and blade/tower loading is provided by OpenFAST to be used as source terms within the momentum equation.

3.15.1 Governing Equations

We begin with a review of the momentum and enthalpy conservation equations within the context of wind farm modeling [CLM+12]. Equation (3.138) shows the Favre-filtered momentum conservation equation (Eq. (3.1)) reproduced here with all the terms required to model a wind farm.

$$\underbrace{\frac{\partial}{\partial t} (\bar{\rho} \tilde{u}_i)}_{\text{I}} + \underbrace{\frac{\partial}{\partial x_j} (\bar{\rho} \tilde{u}_i \tilde{u}_j)}_{\text{II}} = - \underbrace{\frac{\partial p'}{\partial x_j} \delta_{ij}}_{\text{III}} - \underbrace{\frac{\partial \tau_{ij}}{\partial x_j}}_{\text{IV}} - \underbrace{2\bar{\rho} \epsilon_{ijk} \Omega_j u_k}_{\text{V}} + \underbrace{(\bar{\rho} - \rho_o) g_i}_{\text{VI}} + \underbrace{S_i^u}_{\text{VII}} + \underbrace{f_i^T}_{\text{VIII}} \quad (3.138)$$

Term **I** represents the time rate of change of momentum (inertia);

Term **II** represents advection;

Term **III** represents the pressure gradient forces (deviation from hydrostatic and horizontal mean gradient);

Term **IV** represents stresses (both viscous and sub-filter scale (SFS)/Reynolds stresses);

Term **V** describes the influence Coriolis forces due to earth's rotation – see Sec. [Section 3.2.2](#);

Term **VI** describes the effects of buoyancy using the Boussinesq approximation – see [Section 3.2.2](#);

Term **VII** represents the source term used to drive the flow to a horizontal mean velocity at desired height(s) – see [Section 3.2.7](#); and

Term **VIII** is an optional term representing body forces when modeling turbine with actuator disk or line representations – see [Section 3.15.6](#).

In wind energy applications, the energy conservation equation is often written in terms of the Favre-filtered potential temperature, θ , equation, as shown below

$$\frac{\partial}{\partial t} (\bar{\rho} \tilde{\theta}) + \frac{\partial}{\partial x_j} (\bar{\rho} \tilde{u}_j \tilde{\theta}) = - \frac{\partial}{\partial x_j} \hat{q}_j \quad (3.139)$$

where, \hat{q}_j represents the temperature transport due to molecular and SFS turbulence effects. Due to the high Reynolds number associated with ABL flows, the molecular effects are neglected everywhere except near the terrain. Potential temperature is related to absolute temperature by the following equation

$$\theta = T \left(\frac{\bar{p}}{p_o} \right)^{-\left(\frac{R}{c_p}\right)}$$

Under the assumption of ideal gas conditions and constant c_p , the gradients in potential temperature are proportional to the gradients in absolute temperature, i.e.,

$$\left[\frac{\partial T}{\partial t}, \frac{\partial T}{\partial x}, \frac{\partial T}{\partial y} \right] = \left(\frac{\bar{p}}{p_o} \right)^{\left(\frac{R}{c_p}\right)} \left[\frac{\partial \theta}{\partial t}, \frac{\partial \theta}{\partial x}, \frac{\partial \theta}{\partial y} \right]$$

Furthermore, ignoring the pressure and viscous work terms in Eq. (3.12) and assuming constant density (incompressible flow), it can be shown that solving the enthalpy equation is equivalent to solving the potential temperature equation. The enthalpy equation solved in wind energy problems is shown below

$$\frac{\partial}{\partial t} (\bar{\rho} \tilde{T}) + \frac{\partial}{\partial x_j} (\bar{\rho} \tilde{u}_j \tilde{T}) = - \frac{\partial}{\partial x_j} q_j \quad (3.140)$$

It is noted here that the terms \hat{q}_j (Eq. (3.139)) and q_j (Eq. (3.140)) are not equivalent and must be scaled appropriately. User can still provide the appropriate initial and boundary conditions in terms of potential temperature field. Under these assumptions and conditions, the resulting solution can then be interpreted as the variation of potential temperature field in the computational domain.

3.15.2 Turbulence Modeling

LES turbulence closure is provided by the *Subgrid-Scale Kinetic Energy One-Equation LES Model* or the standard *Smagorinsky* model for wind farm applications.

3.15.3 Numerical Discretization & Stabilization

Nalu-Wind provides two discretization approaches

Control Volume Finite Element Method (CVFEM)

Nalu-Wind uses a *dual mesh* approach (see Section 3.3.1) where the *control volumes* are constructed around the nodes of the finite elements within the mesh – see Fig. 3.19. The equations are solved at the *integration points* on the *sub-control surfaces* and/or the *sub-control volumes*.

Edge-Based Vertex Centered Scheme

The edge-based scheme is similar to the finite-volume approach used in SOWFA with the nodes at the *cell center* of the dual mesh.




Fig. 3.19: Schematic of HEX-8 mesh showing the finite elements, nodes, and the associated control volume dual mesh.

The numerical discretization approach is covered in great detail in Section 3.3, the advection and pressure stabilization approaches are documented in Section 3.4 and Section 3.5 respectively. Users are strongly urged to read those sections to gain a thorough understanding of the discretization scheme and its impact on the simulations.

3.15.4 Time stepping scheme

The time stepping method in Nalu-Wind is described in the Fuego theory manual [Tea16] for the backward Euler time discretization. The implementation details of the BDF2 time stepping scheme used in Nalu-Wind is described here. The Navier-Stokes equations are written as

$$\begin{aligned} \mathbf{F}_i(\rho^{n+1}, u_i^{n+1}, P^{n+1}) - \int \frac{\partial \rho u_i}{\partial t} \Big|^{n+1} dV &= 0, \\ \mathbf{F}_i(\rho^{n+1}, u_i^{n+1}, P^{n+1}) - \frac{(\gamma_1 \rho^{n+1} u_i^{n+1} + \gamma_2 \rho^n u_i^n + \gamma_3 \rho^n u_i^{n-1})}{\Delta t} \Delta V &= 0, \end{aligned} \quad (3.141)$$

where

$$\begin{aligned} \mathbf{F}_i(\rho^{n+1} u_i^{n+1}) &= - \int \rho^{n+1} u_i^{n+1} u_j^{n+1} n_j dS + \int \tau_{ij}^{n+1} n_j dS - \int P^{n+1} n_i dS - \int (\rho^{n+1} - \rho_o) g_i dV, \\ &= - \int u_i^{n+1} \dot{m}^{n+1} + \int \tau_{ij}^{n+1} n_j dS - \int P^{n+1} n_i dS - \int (\rho^{n+1} - \rho_o) g_i dV. \end{aligned}$$

and γ_i are factors for BDF2 time discretization scheme (see Section 3.13). As is typical of incompressible flow solvers, the mass flow rate through the sub-control surfaces is tracked independent of the velocity to maintain conservation of mass. The following conventions are used:

$$\begin{aligned} \phi^* &= \text{Predicted value of } \phi \text{ at } n+1 \text{ time step before linear solve} \\ \hat{\phi} = \phi^{**} &= \text{Predicted value of } \phi \text{ at } n+1 \text{ time step after linear solve} \end{aligned}$$

The Newton's method is used along with a linearization procedure to predict a solution to the Navier-Stokes equations at time step $n+1$ as

$$\begin{aligned} \mathbf{A}_{ij} \delta u_j &= \mathbf{F}_i^* - \frac{(\gamma_1 \rho^* u_i^* + \gamma_2 \rho^n u_i^n + \gamma_3 \rho^n u_i^{n-1})}{\Delta t} \Delta V, \\ \text{where } \delta u_j &= u_i^{**} - u_i^*, \\ \mathbf{A}_{ij} &= \left(\frac{\gamma_1 \rho^*}{\Delta t} \Delta V \delta_{ij} - \frac{\partial F_i}{\partial u_j} \Big|^{*} \right), \\ \text{and } \mathbf{F}_i^* &= - \int u_i^* \dot{m}^* + \int \tau_{ij}^* n_j dS - \int P^* n_i dS - \int (\rho^* - \rho_o) g_i dV. \end{aligned} \quad (3.142)$$

After each Newton or *outer* iteration, ϕ^{**} is a better approximation to ϕ^{n+1} compared to ϕ^* . ρ^* and \dot{m}^* are retained constant through each outer iteration. $\mathbf{F}(\rho^* u_i^{**})$ is linear in u_i^* and hence

$$\mathbf{F}_i^* = \frac{\partial F_i}{\partial u_j} \Big|^{*} u_j^* - \int P^* n_i dS - \int (\rho^* - \rho_o) g_i dV \quad (3.143)$$

Applying Eq. (3.143) to Eq. (3.142), we get the linearized momentum predictor equation solved in Nalu-Wind.

$$\begin{aligned} \mathbf{A}_{ij} \delta u_j &= \frac{\partial F_i}{\partial u_j} \Big|^{*} u_j^* - \int P^* n_i dS - \int (\rho^* - \rho_o) g_i dV \\ &\quad - \frac{(\gamma_1 \rho^* u_i^* + \gamma_2 \rho^n u_i^n + \gamma_3 \rho^{n-1} u_i^{n-1})}{\Delta t} \Delta V \\ \mathbf{A}_{ij} \delta u_j &= \left(\frac{\gamma_1 \rho^*}{\Delta t} \Delta V \delta_{ij} - \frac{\partial F_i}{\partial u_j} \Big|^{*} \right) u_j^* - \int P^* n_i dS - \int (\rho^* - \rho_o) g_i dV \\ &\quad - \frac{(\gamma_2 \rho^n u_i^n + \gamma_3 \rho^{n-1} u_i^{n-1})}{\Delta t} \Delta V \\ \mathbf{A}_{ij} \delta u_j &= \mathbf{A}_{ij} u_j^* - \int P^* n_i dS - \int (\rho^* - \rho_o) g_i dV \\ &\quad - \frac{(\gamma_2 \rho^n u_i^n + \gamma_3 \rho^{n-1} u_i^{n-1})}{\Delta t} \Delta V \end{aligned} \quad (3.144)$$

u_i^{**} will not satisfy the continuity equation. A correction step is performed at the end of each outer iteration to make u_i^{**} satisfy the continuity equation as

$$u_i^{n+1} = u_i^{**} - \frac{\tau_3}{\rho} \mathbf{G} \Delta P^{**}$$

where $\Delta P^{**} = P^{**} - P^*$

As described in Section 3.12.1, the continuity equation to be satisfied along with the splitting and stabilization errors is

$$\mathbf{D} \rho u^{**} = b + (\mathbf{L}_1 - \mathbf{D} \tau_3 \mathbf{G}) \Delta P^{**} + (\mathbf{L}_2 - \mathbf{D} \tau_2 \mathbf{G}) P^* \quad (3.145)$$

where b contains any source terms when the velocity field is not divergence free and the other terms are the errors due to pressure stabilization as shown by Domino [Dom06]. The final pressure Poisson equation solved to enforce continuity at each outer iteration is

$$\begin{aligned} u^{n+1} &= u^{**} - \frac{\tau_3}{\rho} \mathbf{G} \Delta P^{**} \\ b + (\mathbf{L}_1 - \mathbf{D} \tau_3 \mathbf{G}) \Delta P^{**} + (\mathbf{L}_2 - \mathbf{D} \tau_2 \mathbf{G}) P^* &= \mathbf{D}(\rho u^{n+1}) = \mathbf{D}(\rho \hat{u}) - \mathbf{D}(\tau_3 \mathbf{G} \Delta P^{**}) \\ b + \mathbf{L}_1 \Delta P^{**} &= \mathbf{D}(\rho \hat{u}) - (\mathbf{L}_2 - \mathbf{D} \tau_2 \mathbf{G}) P^* \\ -\mathbf{L}_1 \Delta P^{**} &= \mathbf{D} \rho \hat{u} + \mathbf{D} \tau_2 \mathbf{G} P^* - \mathbf{L}_2 P^* \\ -\mathbf{L}_1 \Delta P^{**} &= -\mathbf{D} \rho \hat{u} - \mathbf{D} \tau_2 \mathbf{G} P^* + \mathbf{L}_2 P^* + b \end{aligned} \quad (3.146)$$

Thus, the final set of equations solved at each outer iteration is

$$\begin{aligned} \mathbf{A}_{ij} \delta u_j &= \mathbf{A}_{ij} u_j^* - \int P^* n_i dS - \int (\rho^* - \rho_o) g_i dV \\ &\quad - \frac{(\gamma_2 \rho^n u_i^n + \gamma_3 \rho^{n-1} u_i^{n-1})}{\Delta t} \Delta V \\ -\mathbf{L}_1 \Delta P^{**} &= -\mathbf{D} \rho \hat{u} - \mathbf{D} \tau_2 \mathbf{G} P^* + \mathbf{L}_2 P^* + b \\ u_i^{n+1} &= u_i^{**} - \frac{\tau_3}{\rho} \mathbf{G} \Delta P^{**} \end{aligned}$$

Approximations for the Schur complement

Nalu-Wind implements two options for approximating the Schur complement for the split velocity-pressure solution of the incompressible momentum and continuity equation. The two options are:

$$\begin{aligned} \tau &= \tau_1 = \tau_2 = \tau_3 = \Delta t \quad \text{Original implementation} \\ &= (A_{ii})^{-1} \quad \text{Alternate algorithm} \end{aligned}$$

where A_{ii} is the diagonal entry of the momentum linear system. The latter option is similar to the SIMPLE and PIMPLE implementations in OpenFOAM and is used for simulations with RANS and hybrid RANS-LES models with large Courant numbers.

Underrelaxation for momentum and scalar transport

By default, Nalu-Wind applies no underrelaxation during the solution of the Navier-Stokes equations. However, in RANS simulations at large timesteps some underrelaxation might be necessary to restore the diagonal dominance of the transport equations. User has the option to specify underrelaxation through the input files. When underrelaxation is applied, the advection and diffusion contributions to the diagonal term are modified by dividing these terms by

the underrelaxation factor. It must be noted that the underrelaxation is only applied to the advective and viscous contributions in the diagonal term and not the time derivative term.

$$A_{ii} = -\frac{\sum_{i \neq j} A_{ij}}{\omega} + \frac{\gamma_1 \rho \Delta V}{\Delta T}$$

The pressure update can also be underrelaxed by specifying the appropriate relaxation factor in the input file. When this option is activated, the full pressure update, in a given Picard iteration step, is used to project the velocity and mass flow rate and the relaxation is applied to the pressure solution at the end of the Picard iteration.

3.15.5 Initial & Boundary Conditions

This section briefly describes the boundary conditions available in Nalu-Wind for modeling wind farm problems. The terrain and top boundary conditions are described first as they are common to precursor and wind farm simulations.

Initial conditions

Nalu-Wind has the ability to initialize the internal flow fields to uniform conditions for all pressure, velocity, temperature, and TKE (k) in the *input file*. Nalu-Wind also provides a *user function* to add perturbations to the velocity field to trigger turbulence generation during precursor simulations. To specify more complex flow field conditions, a temperature profile with a capping inversion for example, users are referred to pre-processing utilities available in [NaluWindUtils](#) library.

Terrain (Wall) boundary condition

Users are referred to [Section 3.9.3](#) for the treatment of the terrain BC using roughness models. For enthalpy, users can provide a surface heat flux for modeling stratified flows.

Top boundary condition

For problems with minimal streamline curvature near the upper boundary (e.g. nearly flat terrain, negligible turbine blockage), a *symmetry BC* (slip wall) can be when modeling wind farm problems. By default a zero vertical temperature gradient will be imposed for the enthalpy equation when the symmetry boundary condition is used. If a non-zero normal temperature gradient is required to drive the flow to a desired temperature profile, e.g., a capping inversion, then the *abltop BC* can be used. In this case the `user_data` input `normal_temperature_gradient` value will set the normal temperature gradient to value at the top boundary.

For cases with significant terrain features or significant turbine blockage, the *abltop BC* can also be used to achieve an open boundary that allows for both inflows and outflows at the domain top. See the *abltop BC* documentation for details.

Inlet conditions

Time histories of inflow velocity and temperature profiles can be provided as inputs (via I/O transfer) to drive the wind farm simulation with the desired flow conditions. See [Section 4.8.3](#) for more details on this capability. Driving a wind farm simulation using velocity and temperature fields from a mesoscale (WRF) simulation would require an additional pre-processing steps with the [wrftonalu](#) utility.

Outlet conditions

See the description of *open BC* for detailed description of the outlet BC implementation. For wind energy problems, it is necessary to activate the global mass correction as a single value of pressure across the boundary layer is not appropriate in the presence of buoyancy effects. It might also be necessary to fix the reference pressure at an interior node in order to ensure that the Pressure Poisson solver is well conditioned.

3.15.6 Wind Turbine Modeling

Wind turbine rotor and tower aerodynamic effects are modeled using actuator source representations. Compared to resolving the geometry of the turbine, actuator modeling alleviates the need for a complex body-fitted meshes, can relax time step restrictions, and eliminates the need for turbulence modeling at the turbine surfaces. This comes at the expense of a loss of fine-scale detail, for example, the boundary layers of the wind turbine surfaces are not resolved. However, actuator methods well represent wind turbine wakes in the mid to far downstream regions where wake interactions are important.

Actuator methods usually fall within the classes of disks, lines, surface, or some blend between the disk and line (i.e., the swept actuator line). Most commonly, the force over the actuator is computed, and then applied as a body-force source term, f_i (Term **VIII**), to the Favre-filtered momentum equation (Eq. (3.138)).

The body-force term f_i is volumetric and is a force per unit volume. The actuator forces, F'_i , are not volumetric. They exist along lines or on surfaces and are force per unit length or area. Therefore, a projection function, g , is used to project the actuator forces into the fluid volume as volumetric forces. A simple and commonly used projection function is a uniform Gaussian as proposed by Sorensen and Shen [SorensenS02],

$$g(\vec{r}) = \frac{1}{\pi^{3/2}\epsilon^3} e^{-(|\vec{r}|/\epsilon)^2},$$

where \vec{r} is the position vector between the fluid point of interest to a particular point on the actuator, and ϵ is the width of the Gaussian, that determines how diluted the body force become. As an example, for an actuator line extending from $l = 0$ to L , the body force at point (x, y, z) due to the line is given by

$$f_i(x, y, z) = \int_0^L g(\vec{r}(l)) F'_i(l) dl. \quad (3.147)$$

Here, the projection function's position vector is a function of position on the actuator line. The part of the line nearest to the point in the fluid at (x, y, z) has most weight.

The force along an actuator line or over an actuator disk is often computed using blade element theory, where it is convenient to discretize the actuator into a set of elements. For example, with the actuator line, the line is broken into discrete line segments, and the force at the center of each element, F_i^k , is computed. Here, k is the actuator element index. These actuator points are independent of the fluid mesh. The point forces are then projected onto the fluid mesh using the Gaussian projection function, $g(\vec{r})$, as described above. This is convenient because the integral given in Equation (3.147) can become the summation

$$f_i(x, y, z) = \sum_{k=0}^N g(\vec{r}^k) F_i^k. \quad (3.148)$$

This summation well approximates the integral given in Equation (3.147) so long as the ratio of actuator element size to projection function width ϵ does not exceed a certain threshold.

Presently, Nalu-Wind uses an actuator line representation to model the effects of turbine on the flow field; however, the class hierarchy is designed with the potential to add other actuator source terms such as actuator disk, swept actuator line and actuator surface capability in the future. The *ActuatorLineFAST* class couples Nalu-Wind with NREL's OpenFAST for actuator line simulations of wind turbines. OpenFAST is a aero-hydro-servo-elastic tool to model wind turbine developed by the National Renewable Energy Laboratory (NREL). The *ActuatorLineFAST* class allows Nalu-Wind to interface as an inflow module to OpenFAST by supplying the velocity field information.

Nalu-Wind – OpenFAST Coupling Algorithm

A nacelle model is implemented using a Gaussian drag body force. The model implements a drag force in a direction opposite to velocity field at the center of the Gaussian. The width of the Gaussian kernel is determined using the reference area and drag coefficient of the nacelle as shown by Martinez-Tossas. [MartinezT17]

$$\epsilon_d = \sqrt{2 c_d A / \pi}$$

where c_d is the drag coefficient, and A is the reference area. This value of ϵ_d ensures that the momentum thickness of the generated wake is of the right size. The velocity sampled at the center of the Gaussian is corrected to obtain the upstream velocity.

$$\tilde{u}_{i\infty} = \frac{1}{1 - c_d A / (4 \pi \epsilon_d^2)} \tilde{u}_{ic}$$

where \tilde{u}_{ic} is the velocity at the center of the Gaussian and $\tilde{u}_{i\infty}$ is the free-stream velocity used to compute the drag force. The drag body force is then

$$f_d(x, y, z) = \frac{1}{2} c_d A \tilde{u}_{i\infty}^2 \frac{1}{\pi^{3/2} \epsilon_d^3} e^{-|\vec{r}|^2 / \epsilon_d^2}$$

where \vec{r} is the position vector between the fluid point of interest and the center of the Gaussian force.

The actuator line implementation allows for flexible blades that are not necessarily straight (pre-bend and sweep). The current implementation requires a fixed time step when coupled to OpenFAST, but allows the time step in Nalu-Wind to be an integral multiple of the OpenFAST time step. At present, a simple time lagged FSI model is used to interface Nalu-Wind with the turbine model in OpenFAST:

- The velocity at time step at time step n is sampled at the actuator points and sent to OpenFAST,
- OpenFAST advances the turbines up-to the next Nalu-Wind time step $n + 1$,
- The body forces at the actuator points are converted to the source terms of the momentum equation to advance Nalu-Wind to the next time step $n + 1$.

This FSI algorithm is expected to be only first order accurate in time. We are currently working on improving the FSI coupling scheme to be second order accurate in time.

Nalu-Wind – Actuator Disk Model via OpenFAST

An actuator disk model is implemented in Nalu-Wind by using an OpenFAST actuator line to sample the flow and compute the forcing. The actuator line is held stationary which leads to computational savings during execution because there is only 1 search operation in the initial setup.

The forces are gathered at each actuator line point, and the total force at each discrete radial location (r_j where $j \in [1, N_R]$) is computed using `diskTotalForce`.

$$: label : diskTotalForce \mathbf{F}_{total}(r_j) = \sum_{i=1}^{N_B} \mathbf{F}(r_j, \theta_i)$$

where N_B and N_R are the number of blades and number of radial points respectively.

$\mathbf{F}_{total}(r_j)$ is then spread evenly across the original actuator line points and additional ‘swept-points’ that are added in between the actuator lines. The swept-points are always uniformly distributed azimuthally, but the number of swept points can either be non-uniformly or uniformly distributed along the radial direction (left and right images in figure Fig. 3.20). The non-uniform distribution uses the distance between points along the embedded actuator line blades as the arc-length between points in the azimuthal direction. This is the default behavior. If uniform spacing is desired

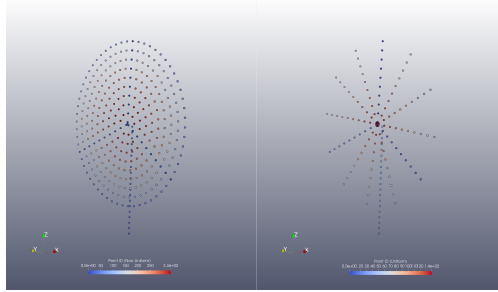


Fig. 3.20: Actuator Disk with non-uniform (left) and uniform (right) sampling in the azimuthal direction.

then `num_swept_pts` must be specified in the input deck. This is the number of points between the actuator lines, so in figure Fig. 3.20 the `num_swept_pts` is 3.

The force that is spread across all the points at a given radius is then calculated as `diskAppliedForce`.

$$: label : diskAppliedForce \mathbf{f}(r_j) = \frac{\mathbf{F}_{total}(r_j)}{N_B * (N_{S,j} + 1)}$$

where $N_{S,j}$ is the number of swept points for a given radius. The index j is used because this value varies between radii when non-uniform sampling is applied.

3.16 Topological Support

The currently supported elements are as follows: hex, tet, pyramid, wedge, quad, and tri. In general, hybrid meshes are fully supported for the edge-based scheme. For CVFEM, hybrid meshes are also supported, however, wedge and pyramid elements are not permitted at exposed open or symmetry boundaries. The remedy to the CVFEM constraint is to simply implement the exposed face gradient operators.

3.17 Adaptivity

Adaptivity is supported through usage of the Percept module. However, this code base has not yet been deployed to the open sector. As such, `ifdef` guards are placed within the code base. A variety of choices exist for the manner by which hanging nodes are removed in a vertex-centered code base.

A typical h-adapted patch of elements is shown in Figure Fig. 3.21. The “hanging nodes” do not have control volumes associated with them. Rather, they are constrained to be a linear combination of the two parent edge nodes. There is no element assembly procedure to compute fluxes for the “hanging sub-faces” associated with the hanging nodes that occur along the parent-child element boundary.

In general, for a vertex-centered scheme, the h-adaptive scheme is driven at the element level. Refinement occurs within the element and the topology of refined elements is the same as the parent element.

Aftosmis [Aft94] describes a vertex-centered finite-volume scheme on unstructured Cartesian meshes. A transitional set of control volumes are formed about the hanging nodes, shown in Figure Fig. 3.22. on unstructured meshes. This approach would require a series of specialized master elements to deal with the different transition possibilities.

Kallinderis [KB89] describes a vertex-centered finite-volume scheme on unstructured quad meshes. Hanging nodes are treated with a constraint condition. The flux construction for a node on a refinement boundary is based on the unrefined parent elements, leading to a non-conservative scheme.




Fig. 3.21: Control volume definition on an h-adapted mesh with hanging nodes. (Four-patch of parent elements with refinement in bottom-right element.)

Kallinderis [KV93] also describes a vertex-centered finite-volume scheme on unstructured tetrahedral meshes. Hanging nodes are removed by splitting the elements on the “unrefined” side of the refinement boundary. Mavriplis [Mav00] uses a similar technique, however, extends it to a general set of heterogeneous elements, shown in Figure Fig. 3.23.

The future deployment of Percept will use the procedure of Mavriplis whereby hanging nodes are removed by neighbor topological changes. A variety of error indicators exists and a prototyped error transport equation approach for the one-equation k^{sgs} model has been tested for classic jet-in-crossflow configurations.

3.17.1 Prolongation and Restriction

Nodal variables are interpolated between levels of the h-adapted mesh hierarchy using the traditional prolongation and restriction operators defined over an element. The prolongation operation is used to compute values for new nodes that arise from element sub-division. The parent element shape functions are used to interpolate values from the parent nodes to the sub-divided nodes.

Prolongation and restriction operators for element variables and face variables are required to maintain mass flow rates that satisfy continuity. When adaptivity takes place, a code option to reconstruct the mass flow rates must be used. Whether or not a Poisson system must be created has been explored. More work is required to understand the nuances associated with prolongation, specifically with respect to possible dispersion errors.

3.18 Code Abstractions

The Nalu-Wind code base is a c++ code-base that significantly leverages the Sierra Toolkit and Trilinos infrastructure. This section is designed to provide a high level overview of the underlying abstractions that the code base exercises. For more detailed code information, the developer is referred to the Trilinos project (github.com). In the sections that follow, only a high level overview is provided.




Fig. 3.22: Control volume definition on an h-adapted mesh with transition control volumes about the hanging nodes. (Four-patch of parent elements with refinement in bottom-right element.)




Fig. 3.23: Control volume definition on a heterogeneous h-adapted mesh with no hanging nodes. (Four-patch of parent elements with refinement in bottom-right element and splitting in adjacent parent elements.)

The Nalu-Wind code base emerged as a small testbed unit test to evaluate the STK infrastructure. Interestingly, the first “algorithm” implementation was a simple L_2 projected nodal gradient. This effort involved reading in a mesh, registering a nodal (vector) field, iterating elements and exposed surfaces to assemble the projected nodal gradient to the nodes of the mesh (in parallel). When evaluating kokkos, this algorithm was also used to learn about the parallel NGP abstraction provided.

3.18.1 Sierra Toolkit Abstractions

Consider a typical mesh that consists of nodes, sides of elements and elements. Such a mesh, when using the Exodus standard, will likely be represented by a collection of “element blocks”, “sidesets” and, possibly, “nodesets”. The definition of the mesh (generated by the user through commercial meshing packages such as pointwise or ICM-CFD) will provide the required spatial definitions of the volume physics and the required boundary conditions.

An element block is a homogeneous collection of elements of the same underlying topology, e.g., HEXAHEDRAL-8. A sideset is a set of exposed element faces on which a boundary condition is to be applied. Finally, a nodeset is a collection of nodes. In general, nodesets are possibly output entities as the code does not exercise enforcing physics or boundary conditions on nodesets. Although Nalu-Wind supports an edge-based scheme, an edge, which is an entity connecting two nodes, is not part of the Exodus standard and must be generated within the STK infrastructure. Therefore, a particular discretization choice may require `stk::mesh::Entity` types of element, face (or side), edge and node.

Once the mesh is read in, a variety of routine operations are generally required. For example, a low-Mach physics equation set may want to be applied to `block_1` while inflow, open, symmetry, periodic and wall boundary conditions can be applied to a variety of sidesets. For example, `surface_1` might be of an “inflow” type. Therefore, the high level set of requirements on a mesh infrastructure might be to allow one to iterate parts of the mesh and, in the end, assemble a quantity to a nodal or elemental field.

Meta and Bulk Data

Meta and Bulk data are simply STK containers. `MetaData` is used to extract parts, extract ownership status, determine the side rank, field declaration, etc. `BulkData` is used to extract buckets, extract upward and downward connectivities and determine node count for a given entity.

Parallel Rules

In STK, elements are locally owned by a single rank. Elements may be ghosted to other parallel ranks through STK custom ghosting. Exposed faces are locally owned by the lower parallel rank. Nodes are also locally owned by the lower parallel rank and can also be shared by all parallel ranks touching them. Edges and internal faces (element:face:element connectivity) have the same rule of locally owned/shared and can also be ghosted. Again, edges and internal faces must be created by existing STK methods should the physics algorithm require them. In Nalu-Wind, the choice of element-based or edge-based is determined within the input file.

Connectivity

In an unstructured mesh, connectivity must be built from the mesh and can not be assumed to follow an assumed “i-j-k” data layout, i.e., structured. In general, one speaks of downward and upward relationships between the underlying entities. For example, if one has a particular element, one might like to extract all of the nodes connected to the element. Likewise, this represents a common operation for faces and edges. Such examples are those in which downward relationships are required. However, one might also have a node and want to extract all of the connected elements to this node (consider some sort of patch recovery algorithm). STK provides the ability to extract such connectivities. In general, full downward and upward connectivities are created.

For example, consider an example in which one has a pointer to an element and wants to extract the nodes of this element. At this point, the developer has not been exposed to abstractions such as buckets, selectors, etc. As such, this is a very high level overview with more details to come in subsequent sections. Therefore, the scope below is to assume that from an element-*k* of a “bucket”, *b[k]* (which is a collection of homogeneous RANK-ed entities) we will extract the nodes of this element using the STK bulk data.

```
// extract element from this bucket
stk::mesh::Entity elem = b[k];

// extract node relationship from bulk data
stk::mesh::Entity const * node_rels = bulkData_.begin_nodes(elem);
int num_nodes = bulkData_.num_nodes(elem);

// iterate nodes
for ( int ni = 0; ni < num_nodes; ++ni ) {
    stk::mesh::Entity node = node_rels[ni];

    // set connected nodes
    connected_nodes[ni] = node;

    // gather some data, e.g., density at state Np1,
    // into a local workset pointer to a std::vector
    p_density[ni] = *stk::mesh::field_data(densityNp1, node );
}
```

Parts

As noted before, a `stk::mesh::Part` is simply an abstraction that describes a set of mesh entities. If one has the name of the part from the mesh data base, one may extract the part. Once the part is in hand, one may iterate the underlying set of entities, walk relations, assemble data, etc.

The following example simply extracts a part for each vector of names that lives in the vector `targetNames` and provides this part to all of the underlying equations that have been created for purposes of nodal field registration. Parts of the mesh that are not included within the `targetNames` vector would not be included in the field registration and, as such, if this missing part was used to extract the data, an error would occur.

```
for ( size_t itarget = 0; itarget < targetNames.size(); ++itarget ) {
    stk::mesh::Part *targetPart = metaData_.get_part(targetNames[itarget]);

    // check for a good part
    if ( NULL == targetPart ) {
        throw std::runtime_error("Trouble with part " + targetNames[itarget]);
    }
    else {
        EquationSystemVector::iterator ii;
        for( ii=equationSystemVector_.begin(); ii!=equationSystemVector_.end(); ++ii )
            (*ii)->register_nodal_fields(targetPart);
    }
}
```

Selectors

In order to arrive at the precise parts of the mesh and entities on which one desires to operate, one needs to “select” what is useful. The STK selector infrastructure provides this.

In the following example, it is desired to obtain a selector that contains all of the parts of interest to a physics algorithm that are locally owned and active.

```
// define the selector; locally owned, the parts I have served up and active
stk::mesh::Selector s_locally_owned_union = metaData_.locally_owned_part()
& stk::mesh::selectUnion(partVec_)
& !(realm_.get_inactive_selector());
```

Buckets

Once a selector is defined (as above) an abstraction to provide access to the type of data can be defined. In STK, the mechanism to iterate entities on the mesh is through the `stk::mesh::bucket` interface. A bucket is a homogeneous collection of `stk::mesh::Entity`.

In the below example, the selector is used to define the bucket of entities that are provided to the developer.

```
// given the defined selector, extract the buckets of type ``element``
stk::mesh::BucketVector const& elem_buckets
= bulkData_.get_buckets( stk::topology::ELEMENT_RANK,
                        s_locally_owned_union );

// loop over the vector of buckets
for ( stk::mesh::BucketVector::const_iterator ib = elem_buckets.begin();
      ib != elem_buckets.end() ; ++ib ) {
    stk::mesh::Bucket & b = **ib ;
    const stk::mesh::Bucket::size_type length = b.size();

    // extract master element (homogeneous over buckets)
    MasterElement *meSCS = sierra::nalu::get_surface_master_element(b.topology());

    for ( stk::mesh::Bucket::size_type k = 0 ; k < length ; ++k ) {

        // extract element from this bucket
        stk::mesh::Entity elem = b[k];

        // etc...
    }
}
```

The look-and-feel for nodes, edges, face/sides is the same, e.g.,

- for nodes:

```
// given the defined selector, extract the buckets of type ``node``
stk::mesh::BucketVector const& node_buckets
= bulkData_.get_buckets( stk::topology::NODE_RANK,
                        s_locally_owned_union );

// loop over the vector of buckets
```

- for edges:

```
// given the defined selector, extract the buckets of type ``edge``
stk::mesh::BucketVector const& edge_buckets
= bulkData_.get_buckets( stk::topology::EDGE_RANK,
                        s_locally_owned_union );
```

(continues on next page)

(continued from previous page)

```
// loop over the vector of buckets
```

- for faces/sides:

```
// given the defined selector, extract the buckets of type ``face/side``
stk::mesh::BucketVector const& face_buckets
    = bulkData_.get_buckets( metaData_.side_rank(),
                           s_locally_owned_union );

// loop over the vector of buckets
```

Field Data Registration

Given a part, we would like to declare the field and put the field on the part of interest. The developer can register fields of type elemental, nodal, face and edge of desired size.

- nodal field registration:

```
void
LowMachEquationSystem::register_nodal_fields(
    stk::mesh::Part *part)
{
    // how many states? BDF2 requires Np1, N and Nm1
    const int numStates = realm_.number_of_states();

    // declare it
    density_
        = &(metaData_.declare_field<ScalarFieldType>(stk::topology::NODE_RANK,
                                                       "density", numStates));

    // put it on this part
    stk::mesh::put_field(*density_, *part);
}
```

- edge field registration:

```
void
LowMachEquationSystem::register_edge_fields(
    stk::mesh::Part *part)
{
    const int nDim = metaData_.spatial_dimension();
    edgeAreaVec_
        = &(metaData_.declare_field<VectorFieldType>(stk::topology::EDGE_RANK,
                                                       "edge_area_vector"));
    stk::mesh::put_field(*edgeAreaVec_, *part, nDim);
}
```

- side/face field registration:

```
void
MomentumEquationSystem::register_wall_bc(
    stk::mesh::Part *part,
    const stk::topology &theTopo,
    const WallBoundaryConditionData &wallBCData)
```

(continues on next page)

(continued from previous page)

```

{
    // Dirichlet or wall function bc
    if ( wallFunctionApproach ) {
        stk::topology::rank_t sideRank
            = static_cast<stk::topology::rank_t>(metaData_.side_rank());
        GenericFieldType *wallFrictionVelocityBip
            = &(metaData_.declare_field<GenericFieldType>
                (sideRank, "wall_friction_velocity_bip"));
        stk::mesh::put_field(*wallFrictionVelocityBip, *part, numIp);
    }
}

```

Field Data Access

Once we have the field registered and put on a part of the mesh, we can extract the field data anytime that we have the entity in hand. In the example below, we extract nodal field data and load a workset field.

To obtain a pointer for a field that was put on a node, edge face/side or element field, the string name used for declaration is used in addition to the field template type,

```

VectorFieldType *velocityRTM
    = metaData_.get_field<VectorFieldType>(stk::topology::NODE_RANK,
                                           "velocity");

ScalarFieldType *density
    = metaData_.get_field<ScalarFieldType>(stk::topology::NODE_RANK,
                                           "density");

VectorFieldType *edgeAreaVec
    = metaData_.get_field<VectorFieldType>(stk::topology::EDGE_RANK,
                                           "edge_area_vector");

GenericFieldType *massFlowRate
    = metaData_.get_field<GenericFieldType>(stk::topology::ELEMENT_RANK,
                                           "mass_flow_rate_scs");

GenericFieldType *wallFrictionVelocityBip_
    = metaData_.get_field<GenericFieldType>(metaData_.side_rank(),
                                           "wall_friction_velocity_bip");

```

State

For fields that require state, the field should have been declared with the proper number of states (see field declaration section). Once the field pointer is in hand, the specific field with state is easily extracted,

```

ScalarFieldType *density
    = metaData_.get_field<ScalarFieldType>(stk::topology::NODE_RANK,
                                           "density");

densityNm1_ = &(density->field_of_state(stk::mesh::StateNM1));
densityN_   = &(density->field_of_state(stk::mesh::StateN));
densityNp1_ = &(density->field_of_state(stk::mesh::StateNP1));

```

With the field pointer already in hand, obtaining the particular data is field the field data method.

- nodal field data access:

```
// gather some data (density at state Np1) into a local workset pointer
p_density[ni] = *stk::mesh::field_data(densityNp1, node );
```

- **edge field data access:** (from an edge bucket loop with the same selector as defined above)

```
stk::mesh::BucketVector const& edge_buckets
= bulkData_.get_buckets( stk::topology::EDGE_RANK, s_locally_owned_union );
for ( stk::mesh::BucketVector::const_iterator ib = edge_buckets.begin();
      ib != edge_buckets.end() ; ++ib ) {
    stk::mesh::Bucket & b = **ib ;
    const stk::mesh::Bucket::size_type length = b.size();

    // pointer to edge area vector and mdot (all of the buckets)
    const double * av = stk::mesh::field_data(*edgeAreaVec_, b);
    const double * mdot = stk::mesh::field_data(*massFlowRate_, b);

    for ( stk::mesh::Bucket::size_type k = 0 ; k < length ; ++k ) {
        // copy edge area vector to a pointer
        for ( int j = 0; j < nDim; ++j )
            p_areaVec[j] = av[k*nDim+j];

        // save off mass flow rate for this edge
        const double tmdot = mdot[k];
    }
}
```

3.18.2 High Level Nalu-Wind Abstractions

Realm

A realm holds a particular physics set, e.g., low-Mach fluids. Realms are coupled loosely through a transfer operation. For example, one might have a turbulent fluids realm, a thermal heat conduction realm and a PMR realm. The realm also holds a BulkData and MetaData since a realm requires fields and parts to solve the desired physics set.

EquationSystem

An equation system holds the set of PDEs of interest. As Nalu-Wind uses a pressure projection scheme with split PDE systems, the pre-defined systems are, LowMach, MixtureFraction, Enthalpy, TurbKineticEnergy, etc. New monolithic equation system can be easily created and plugged into the set of all equation systems.

In general, the creation of each equation system is of arbitrary order, however, in some cases fields required for MixtureFraction, e.g., `mass_flow_rate` might have only been registered on the low-Mach equation system. As such, if MixtureFraction is created before LowMachEOS, an error might be noted. This can be easily resolved by cleaning the code base such that each equation system is “autonomous”.

Each equation system has a set of virtual methods expected to be implemented. These include, however, are not limited to registration of nodal fields, edge fields, boundary conditions of fixed type, e.g., wall, inflow, symmetry, etc.

NALU-WIND - VERIFICATION MANUAL

Nalu-Wind represents a generalized unstructured, massively parallel, variable density turbulent flow capability designed for energy applications. This code base began as an effort to prototype Sierra Toolkit, [EWS+10], usage along with direct parallel matrix assembly to the Trilinos, [HBH+03], Epetra and Tpetra data structure. However, the simulation tool has evolved as a tool to support a variety of research projects germane to the energy sector including wind aerodynamic prediction and traditional gas-phase combustion applications.

4.1 Introduction

The methodology used to evaluate the accuracy of each proposed scheme will be the method of manufactured solutions. The objective of code verification is to reveal coding mistakes that affect the order of accuracy and to determine if the governing discretized equations are being solved correctly. Quite often, the process of verification reveals algorithmic issues that would otherwise remain unknown.

In practice, a variety of comparison techniques exist for verification. For example, benchmark and code-to-code comparison are not considered rigorous due to the errors that exist in other code solutions, such as from discretization and iteration. Analytic solutions and the method of manufactured solutions remain the most powerful methods for code verification, since they provide a means to obtain quantitative error estimations in space and time.

Roache has made the distinction between code verification and calculation verification, where calculation verification involves grid refinement required for every problem solution to assess the magnitude, not order, of the discretization error. Discretization error, distinguished from modeling and iteration errors, is defined as the difference between the exact solution to the continuum governing equations and the solution to the algebraic systems representation due to discretization of the continuum equations. The order of accuracy can be determined by comparing the discretization error on successively refined grids. Thus, it is desirable to have an exact solution for comparison to determine the discretization errors.

4.2 2D Unsteady Uniform Property: Convecting Decaying Taylor Vortex

Verification of first-order and second-order temporal accuracy for the CVFEM and EBVC formulation in Nalu-Wind is performed using the method of manufactured solution (MMS) technique. For the unsteady isothermal, uniform laminar physics set, the exact solution of the convecting, decaying Taylor vortex is used.

$$u = u_o - \cos(\pi(x - u_o t)) \sin(\pi(y - v_o t)) e^{-2.0\omega t} \quad (4.1)$$

$$v = v_o + \sin(\pi(x - u_o t)) \cos(\pi(y - v_o t)) e^{-2.0\omega t} \quad (4.2)$$

$$p = -\frac{p_o}{4} (\cos(2\pi(x - u_o t)) + \cos(2\pi(y - v_o t))) e^{-4\omega t} \quad (4.3)$$

In this study, the constants u_o , v_o , and p_o are all assigned values of 1.0, and the viscosity μ is set to a constant value of 0.001. The value of ω is $\pi^2\mu$. This particular viscosity value results in a maximum cell reynolds number of twenty.

4.2.1 Temporal Order Of Accuracy Results

The temporal order of accuracy for the first order backward Euler and second order BDF2 are outlined in Figure Fig. 4.1 and Figure Fig. 4.2. Each of these simulations used a hybrid factor of zero to ensure pure second order central usage. A fixed Courant number of two was used for each of the three meshes (100x100, 200x200 and 400x400). The simulation was run out to 0.2 seconds and L_2 error norms were computed. The standard fourth order pressure stabilization scheme with time step scaling is used. This scheme is also known as the standard incremental pressure, approximate pressure projection scheme.

Two other pressure projection schemes have been evaluated in this study. Each represent a simplification of the standard pressure projection scheme. Figure Fig. 4.3 outlines three projection schemes: the first is when the projected nodal gradient appearing in the fourth-order pressure stabilization is lagged while the second is the classic pressure-free pressure approximate projection scheme with second order pressure stabilization. The third is the baseline fourth-order incremental pressure projection scheme. The error plots demonstrate that lagging the projected nodal gradient for pressure retains second order accuracy. However, as expected the pressure free pressure projection scheme is confirmed to be first order accurate given the first order splitting error noted in this fully implicit momentum solve.

The Steady Taylor Vortex will be used to verify the spatial accuracy for the full set of advection operators supported in Nalu-Wind.

4.3 Higher Order 2D Steady Uniform Property: Taylor Vortex

A higher order unstructured CVFEM method has been developed by Domino [Dom14]. A 2D structured mesh study demonstrating second order time and third order in space scheme has been demonstrated. The below work has emphasis on unstructured meshes.

4.3.1 Source Term Quadrature

Higher order accuracy is only demonstrated on solutions with source terms when a fully integrated approach is used. Lumping the source term evaluation is a second order error and is fully noted in the MMS study (not shown).

4.3.2 Projected nodal gradients

Results show that one must use design order projected nodal gradients. Figure Fig. 4.4 demonstrates a code verification result for a steady thermal manufactured solution comparing lumped and consistent mass matrix approaches for the projected nodal gradient on a quadratic tqquad mesh. In the lumped approach, a simple explicit algorithm is processed while for the consistent approach, a simple mass matrix inversion equation must be solved. The lumped approach is first order while the consistent approach retains the expected second order as the projected nodal gradient is expected to be order P . Both Dirichlet and periodic domains display the same order of convergence.

4.3.3 Momentum and Pressure

The steady taylor vortex exact solution was run on a quadratic tqquad mesh. Figure Fig. 4.5 demonstrates the order of accuracy for projected nodal gradients (pressure) and the velocity field (x-component). Second order accuracy for the projected nodal gradient (pressure) and third order for the velocity field is realized when the consistent mass matrix approach is used for the projected nodal pressure gradient. Note that this term is used in the pressure stabilization approach. However, order of convergence for the projected nodal pressure gradient and velocity field is compromised




Fig. 4.1: Error norms as a function of timestep size for the u and v component of velocity using fourth order pressure stabilization with timestep scaling, backward Euler




Fig. 4.2: Error norms as a function of timestep size for the u and v component of velocity using fourth order pressure stabilization with timestep scaling, BDF2




Fig. 4.3: Error norms as a function of timestep size for the u and v component of velocity using the lagged projected nodal pressure gradient and pressure-free pressure projection scheme; all with timestep scaling, BDF2




Fig. 4.4: Error norms as a function of mesh size for a CMM and LMM projected nodal gradient on a quadratic quad mesh.

when the lumped mass matrix approach is used for the pressure stabilization term. Note that both approaches use the fully integrated pressure gradient term in the momentum equation (i.e., $\int p n_i dS$). Therefore, the reduced order of integration for the projected nodal pressure gradient has consequence on the velocity field order of convergence.

Again, dirichlet (inflow) and periodic domains display the same order of convergence.




Fig. 4.5: Error norms as a function of mesh size for the Steady Taylor Vortex momentum and pressure gradient field.

4.4 3D Steady Non-isothermal with Buoyancy

Building from the basic functional form of the Taylor Vortex, a non-isothermal solution (momentum, pressure and static enthalpy) is manufactured as follows:

$$\begin{aligned}
 u &= -u_o \cos(a\pi x) \sin(a\pi y) \sin(a\pi z) \\
 v &= +v_o \sin(a\pi x) \cos(a\pi y) \sin(a\pi z) \\
 w &= -w_o \sin(a\pi x) \sin(a\pi y) \cos(a\pi z) \\
 p &= -\frac{p_o}{4} (\cos(2a\pi x) + \cos(2a\pi y) + \cos(2a\pi z)) \\
 h &= +h_o \cos(a_h \pi x) \cos(a_h \pi y) \cos(a_h \pi z)
 \end{aligned} \tag{4.4}$$

The equation of state is simply the ideal gas law,

$$\rho = \frac{P^{ref} M}{RT} \tag{4.5}$$

The simulation is run on a three-dimensional domain ranging from -0.05:+0.05 with constants $a, a_h, M, R, C_p, P^{ref}, T_{ref}, Pr, \mu$ equal to (20, 10, 30, 10, 0.01, 100, 300, 0.8, 0.00125), respectively.

At reference conditions, the density is unity. The effects of buoyancy are also provided by an arbitrary gravity vector of magnitude of approximately ten, $g_i = (-5, 6, 7)^T$. On this domain, the enthalpy ranges from zero to unity. Given the reference values, the temperature field ranges from 300K to 400K which is designed to mimic a current LES non-isothermal validation suite.

Edge- and element-based discretization (P=1) demonstrate second order convergence in the L_2 norm for u, v, w and temperature. This test is captured within the variableDensityMMS regression test suite.

4.5 3D Steady Non-uniform with Buoyancy

Building from the basic functional form of the Taylor Vortex, a non-uniform solution (momentum, pressure and mixture fraction) is manufactured as follows:

$$\begin{aligned} u &= -u_o \cos(a\pi x) \sin(a\pi y) \sin(a\pi z) \\ v &= +v_o \sin(a\pi x) \cos(a\pi y) \sin(a\pi z) \\ w &= -w_o \sin(a\pi x) \sin(a\pi y) \cos(a\pi z) \\ p &= -\frac{p_o}{4} (\cos(2a\pi x) + \cos(2a\pi y) + \cos(2a\pi z)) \\ z &= +z_o \cos(a_z \pi x) \cos(a_z \pi y) \cos(a_z \pi z) \end{aligned} \tag{4.6}$$

The equation of state is simply the standard inverse mixture fraction property expression for density,

$$\rho = \frac{1}{\frac{z}{\rho^P} + \frac{1-z}{\rho^S}} \tag{4.7}$$

The simulation is run on a three-dimensional domain ranging from -0.05:+0.05 with constants $a, a_z, \rho^P, \rho^S, Sc, \mu$ equal to (20, 10, 0.1, 1.0, 0.8, 0.001), respectively.

At reference conditions, the density is that of the primary condition (0.1). The effects of buoyancy are also provided by an arbitrary gravity vector of magnitude of approximately ten, $g_i = (-5, 6, 7)^T$. On this domain, the mixture fraction ranges from zero to unity. This test case is designed to support the helium plume DNS study with primary and secondary density values of helium and air, respectively.

Edge- and element-based discretization (P=1) demonstrate second order convergence in the L_2 norm for u, v, w and mixture fraction. This test is captured within the variableDensityMMS regression test suite.

4.6 2D Steady Laplace Operator

The evaluation of the low-Mach Laplace (or diffusion operator) is of great interest to the core supported application space. Although the application space for Nalu-Wind is characterized by a highly turbulent flow, the usage of an approximate pressure projection scheme always makes the chosen Laplace form important. Although the element-based scheme is expected to be accurate, it can be problematic on high aspect ratio meshes as element-based schemes are not guaranteed to be monotonic for aspect ratios as low as $\sqrt{2}$ for FEM-based schemes and $\sqrt{3}$ for CVEM-based approaches (both when using standard Gauss point locations). Conversely, while the edge-based operator is accurate on high aspect ratio meshes, it suffers on skewed meshes due to both quadrature error and the inclusion of a non orthogonal correction (NOC).

In order to assess the accuracy of the Laplace operator, a the two-dimensional MMS temperature solution is used. The functional temperature field takes on the following form:

$$T = \frac{\lambda}{4} (\cos(2a\pi x) + \cos(2a\pi y)). \tag{4.8}$$

The above manufactured solution is run on three meshes of domain size of 1x1. The domain was first meshed as a triangular mesh and then converted to a tquad4 mesh. Therefore, non orthogonal correction (NOC) effects are expected for the edge-based scheme. In this study, both λ and a are unity. Either periodic or Dirichlet conditions are used for boundary conditions.

A brief overview of the diffusion operator tested is now provided. For more details, consult the theory manual. The general diffusion kernel is as follows:

$$- \int \Gamma \frac{\partial \phi}{\partial x_j} A_j. \quad (4.9)$$

The choice of the gradient operator at the integration point is a function of the underlying method. For CVFEM, the gradient operator is provided by the standard shape function derivatives,

$$\frac{\partial \phi_{ip}}{\partial x_j} = \sum \frac{\partial N_{j,k}^{ip}}{\partial x_j} \phi_k. \quad (4.10)$$

For the edge-based scheme, a blending of an orthogonal gradient along the edge and a NOC is employed,

$$\frac{\partial \phi_{ip}}{\partial x_j} = \bar{G}_j \phi + [(\phi_R - \phi_L) - \bar{G}_l \phi dx_l] \frac{A_j}{A_k dx_k}. \quad (4.11)$$

In the above equation, $\bar{G}_j \phi$ is a projected nodal gradient. The general equation for this quantity is

$$\int w_I \bar{G}_j \phi dV = \int w_i \frac{\partial \phi}{\partial x_j} dV. \quad (4.12)$$

Possible forms of this include either lumped or consistent mass (the later requires a global equation solve) with either the full CVFEM stencil or the edge-based (reduced) stencil. The above equation can even be applied within the element itself for a simple, local integration step that provides a piecewise constant gradient over the element.

The simulation study is run with the following diffusion operators: 1) the standard CVFEM operator, 2) the edge-based operator with CVFEM projected nodal gradients (NOC), 3) the edge-based operator with edge-based projected nodal gradients (NOC), 4) the edge-based operator without NOC correction, 5) the CVFEM operator with shifted integration points to the edge, and, lastly, 6) a mixed edge/element scheme in which the orthogonal diffusion operator is edge-based while the NOC terms are based on the elemental CVFEM gradient (either evaluated at the given integration point or integrated over the element for a piecewise constant form).

The last operator is interesting in that it represents a candidate operator for the CVFEM pressure Poisson system when high aspect ratio meshes are used. Figure Fig. 4.6 outlines the convergence of the five above operators; shown are all of the standard norms (∞ , 1 and 2) for the R0, R1 and R2 mesh refinements. The results in the left side of the figure indicate that the edge-based scheme with NOC retains second-order convergence for all norms when the more accurate CVFEM projected nodal gradient is used (lumped only tested given its good results). Convergence is degraded with the edge-based scheme when NOC terms are either neglected or use the reduced edge-based projected nodal gradient. The CVFEM-based methods are second order accurate in the L_1 and L_2 norms, however, questionable results are noted in the L_∞ norm for all methods that include any shape function derivative for local or elemental piecewise constant gradient operators. Shifting the Gauss points from the standard subcontrol surface to the edges of the element (while still using shape function derivatives) is only problematic in the L_∞ norm (just as the standard CVFEM approach). The use of the mixed-approach with a piecewise constant gradient over the element demonstrates the same behavior as when using the integration point CVFEM gradient operator. Figure Fig. 4.7 outlines two more refinement meshes for the CVFEM operator (R3 and R4). Results indicate that the L_∞ norm is approaching second order accuracy.

An inspection of the magnitude of error between the exact and computed temperature for the R3 mesh is shown in Figure Fig. 4.8. Results show that the CVFEM error is highest at the corner mesh nodes that form a reduced stencil. The edge-based scheme shows increased error at the higher aspect ratio dual mesh.




Fig. 4.6: Error norms for tqquad4 refinement study. R0, R1, and R2 refinement.




Fig. 4.7: Error norms for tqquad4 refinement study. R0, R1, R2, R3, R4, and R4 refinementError for CVFEM.




Fig. 4.8: Magnitude of the L_∞ temperature norm comparing the edge-based CVFEM (NOC) and standard CVFEM operators on the R3 mesh.

4.7 3D Steady Laplace Operator with Nonconformal Interface

A three dimensional element-based verification study is provided to evaluate the DG-based CVFEM approach.

$$T = \frac{\lambda}{4}(\cos(2a\pi x) + \cos(2a\pi y) + \cos(2a\pi z)). \quad (4.13)$$

Figure Fig. 4.9 represents the MMS field for temperature. The simulation study includes uniform refinement of a first- and second-order CVFEM basis. Both temperature field and projected nodal gradient norms are of interest.

Figure Fig. 4.10 outlines the linear and quadratic basis. For P1, the CVFEM temperature field predicts between second and first order while for P2, third order is recovered. When using a consistent mass matrix for the projected nodal gradient, second order is noted, see Figure Fig. 4.11.

dof	L_∞	L1	L2
temperature	3.33067e-16	2.30077e-17	4.68103e-17
dTdx	4.13225e-13	9.06848e-15	1.98249e-14
dTdy	4.15668e-13	1.11256e-14	2.15065e-14
dTdz	4.31211e-13	9.60785e-15	1.97517e-14

Given the order of accuracy results for the P1 implementation, a linear patch test was run. The temperature solution was simply, $T(x, y, z) = x + y + z$; all analytical temperature gradients are unity. Table Table 4.7 demonstrates the successful patch test results for a P1 CVFEM implementation.

4.8 Precursor-based Simulations

In the field of turbulent flow modeling and simulation, often times simulations may require sophisticated boundary conditions that can only be obtained from previously run high-fidelity simulations. For example, consider a typical turbulent jet simulation in which the experimental inlet condition was preceeded by a turbulent pipe entrance region. Furthermore, in most cases the ability to adequately predict the developing jet flow regime may be highly sensitive to proper inlet conditions. Figure Fig. 4.12 and Figure Fig. 4.13 outline a process in which a high fidelity large-eddy simulation of a periodic pipe was used to determine a representative inlet condition for a turbulent round jet. Specifically, a precursor pipe flow simulation is run with velocity provided to an output file. This output file serves as the inlet velocity profile for the subsequent simulation.

In the above use case, as with most general simulation studies, the mesh resolution for the precursor simulation may be different from the subsequent simulation. Moreover, the time scale for the precursor simulation may be much shorter than the subsequent simulation. Finally, the data required for the subsequent simulation will likely be at different time steps unless an overly restrictive rule is enforced, i.e., a fixed timestep for each simulation.

In order to support such use cases, extensive usage of the the Sierra Toolkit infrastructure is expected, most notably within the IO and Transfer modules. The IO module can be used to interpolate the precursor simulation boundary data to the appropriate time required by the subsequent simulation. Specifically, the IO module linearly interpolates between the closest data interval in the precursor data set. A recycling offset factor is included within the IO interface that allows for the cycling of data over the full time scale of interest within the subsequent simulation. For typical statistically stationary turbulent flows, this is useful to ensure proper statistics are captured in subsequent runs.

After the transient data set from the precursor simulation is interpolated to the proper time, the data is spatially interpolated and transferred to the subsequent simulation mesh using the STK Transfer module. Efficient coarse parallel searches (point/bounding box) provide the list of candidate owning elements on which the fine-scale search operates to determine the best search candidate. The order of spatial interpolation depends on the activated numerical discretization. Therefore, by combining the two STK modules, the end use case to support data transfers of boundary data is supported.




Fig. 4.9: MMS temperature field for nonconformal algorithm.




Fig. 4.10: MMS order of accuracy for nonconformal algorithm. Temperature norms for P1 and P2 elements.




Fig. 4.11: MMS order of accuracy for nonconformal algorithm. Projected nodal gradient norms for P1 and P2 elements.




Fig. 4.12: Precursor periodic pipe flow large-eddy simulation that will serve as the inlet boundary condition for a subsequent turbulent jet simulation.




Fig. 4.13: Subsequent turbulent jet simulation using the precursor data obtained by a periodic pipe flow.

As noted, there are many other use cases in addition to the overviewed turbulent jet simulation that require such temporal/spatial interpolation capabilities. For example, in typical wind farm simulation applications, a proper atmospheric boundary layer (ABL) configuration is required to capture a given energy state of the boundary layer. In this case, a periodic precursor ABL is run with the intent of providing the inlet condition to the subsequent wind farm domain. As with the previous description, the infrastructure requirements remain the same.

Finally, the general creation of an “input_output” region can be useful in validation cases where data are provided at a subset of the overall simulation domain. Such is the case in PIV and PLIF experimental data sets. Although the temporal interpolation is not required, the transfer of this data at high time step frequency is useful for post-processing.

In this verification section, a unit test approach will be referenced that is captured within the STK module test suite. This foundational test coverage provides confidence in the underlying IO and parallel search/interpolation processes. In addition to briefly describing the infrastructure testing, application tests are provided as further evidence of correctness. The application test first is based on the convecting Taylor vortex verification case while the second is the ABL precursor application space demonstration.

4.8.1 Infrastructure Unit Test

As noted above, the Nalu-Wind application code leverages the STK unit tests within the IO and transfer modules. Interested parties may peruse the STK product under a cloned Trilinos cloned project, i.e., Trilinos/packages/stk/stk_doc_test. Under the STK product, a variety of search, transfer and input/output tests exist. For example, interpolation in time using the IO infrastructure is captured in addition to a variety of search and transfer use cases.

4.8.2 Application Verification Test; Convecting Taylor Vortex

Although the foundational infrastructure tests are useful, the application must adequately interface the IO and Transfer modules to support the end use case. In this section, two tests will be demonstrated that illustrate the precursor/subsequent simulation use case.

The first test considered will be the convecting Taylor vortex. In this configuration, a very fine mesh simulation is run with boundary conditions specified in the input file to be of type, “convecting_taylor_vortex”. This specifies the analytical function for the x-component of velocity as provided in Equation (4.1). The simulation is run while providing output to a Realm of type “input_output” using a transfer objective, “input_output”. The transient data is then used for a series of mesh refinement studies. The viscosity is set at 0.001 while the domain is 1x1. In this study, the edge-based scheme is activated, however, the precursor interpolation methodology is not sensitive to the underlying numerical method.

In Figure Fig. 4.14, a plot between the analytical x-component of velocity and a nodal query of the outputted velocity component is provided. Although not immediately apparent, the values are exactly the same. This finding confirms that the data set output is consistent with the nodal exact value.

With the precursor data base containing the full transient data, a refinement study can be accomplished to determine numerical errors. Although the full machinery for temporal and spatial interpolation is active, the data requirement at the coarse simulations are represented as the subsets of the full data - both in space and time. As such, no numerical degradation of second-order accuracy is expected. The subsequent simulations are run with an “external_data” transfer objective and a Realm of type, “external_data_provider”.

In Figure Fig. 4.15, a plot of L_2 norms of the x-component of velocity are shown for the subsequent set of simulations that use the precursor data. Results of this study verify both the second-order temporal accuracy of the underlying numerical scheme and the process of using both space and time interpolation.




Fig. 4.14: Temporal plot of the exact x-velocity component and precursor output.




Fig. 4.15: Temporal accuracy plot of the x-velocity component norms using the precursor data.

4.8.3 Application Verification Test; ABL Precursor/Subsequent

The second, and final application test is an ABL-based simulation that first runs a precursor periodic solution in order to capture an appropriate ABL specification. The boundary data saved from the precursor simulation are then used as an inflow boundary condition for the subsequent ABL simulation. As the precursor is run for a smaller time frame than the subsequent simulation, the usage of data cycling is active. This full integration test is captured within the regression test suite. The simulation is described as a non-isothermal turbulent flow.

In Figure Fig. 4.16, the transient recycling of the ABL thermal inflow boundary condition is captured at an arbitrary nodal location very near the wall boundary condition. The subsequent simulation reads the precursor data set for time zero seconds until 3000 seconds at which time it recycles the inlet condition back to the initial precursor simulation time, i.e., zero seconds. An interesting note in this study is the fact that the precursor periodic simulation, which was run at the same Courant number, was using time steps approximately three times greater than the subsequent inflow/open configuration.

In Figure Fig. 4.17, (left) the subsequent simulation inflow temperature field and full profile over the full domain is captured at approximately 4620 seconds. On the right of the figure, the temperature boundary condition data that originated from the precursor simulation, which was read into the subsequent “external_field_provider” Realm, is shown (again at approximately 4620 seconds).

4.9 Boussinesq Verification

4.9.1 Unit tests

Unit-level verification was performed for the Boussinesq body force term (3.9) with a nodal source appropriate to the edge-based scheme (MomentumBoussinesqSrcNodeSuppAlg.single_value) as well as a separate unit test for the element-based “consolidated” Boussinesq source term (MomentumKernelHex8Mesh.buoyancy_boussinesq). Proper volume integration with different element topologies is also tested (the “volume integration” tests in the MasterElement and HOMasterElement test cases).

4.9.2 Stratified MMS

A convergence study using the method of manufactured solutions (MMS) was also performed to assess the integration of the source term into the governing equations. An initial condition of a Taylor-Green vortex for velocity, a zero-gradient pressure field, and a linear enthalpy profile in the z-direction are imposed.

$$\begin{aligned}u &= -\frac{1}{2}\cos(2\pi x)\sin(2\pi y)\sin(2\pi z) \\v &= \sin(2\pi x)\cos(2\pi y)\sin(2\pi z) \\w &= -\frac{1}{2}\sin(2\pi x)\sin(2\pi y)\cos(2\pi z) \\p &= 0 \\h &= z.\end{aligned}\tag{4.14}$$

The simulation is run on a three-dimensional domain ranging from $-1/2$ to $1/2$ with reference density, reference temperature and the thermal expansion coefficient to equal to 1, 300, and 1, respectively. β is much larger than typical ($1/T_{\text{ref}}$) so that the buoyancy term is a significant term in the MMS in this configuration.

The Boussinesq buoyancy model uses a gravity vector of magnitude of ten in the z-direction opposing the enthalpy gradient, $g_i = (0, 0, -10)^T$. The temperature for this test ranges between 250K and 350K. The test case was run with a regular hexahedral mesh, using the edge-based vertex centered finite volume scheme. Each case was run with a fixed maximum Courant number of 0.8 relative to the specified solution.




Fig. 4.16: Transient recycling of the temperature inflow boundary condition for the subsequent ABL simulation. After 3000 seconds, the inflow boundary condition is recycled from the beginning of the precursor simulation.

Time: 4619.539389




Fig. 4.17: Subsequent simulation showing the full temperature domain (left) and on the precursor inflow temperature boundary condition field obtained from the perspective of the subsequent “external_field_provider” Realm (right).

Table 4.1: Error in x-component of velocity

h	L_∞	L1	L2	Order
1/32	8.91e-3	1.12e-3	1.77e-3	NA
1/64	2.03e-3	3.04e-4	4.27e-4	2.05
1/128	4.65e-4	7.64e-5	1.05e-4	2.03

Table 4.2: Error in y-component of velocity

h	L_∞	L1	L2	Order
1/32	1.78e-2	2.31e-3	3.47e-3	NA
1/64	4.18e-3	5.92e-4	8.23e-4	2.06
1/128	9.70e-4	1.50e-4	2.02e-4	2.03

Table 4.3: Error in z-component of velocity

h	L_∞	L1	L2	Order
1/32	8.68e-2	1.17e-3	1.73e-3	NA
1/64	2.00e-3	2.99e-4	4.22e-4	2.04
1/128	4.64e-4	7.63e-5	1.05e-4	2.00

Table 4.4: Error in temperature

h	L_∞	L1	L2	Order
1/32	1.09e-2	1.46e-3	2.10e-3	NA
1/64	2.06e-3	3.13e-4	4.19e-4	2.32
1/128	4.18e-4	7.54e-5	1.00e-4	2.06

This test is added to Nalu-Wind’s nightly test suite, testing that the convergence rate between the 1/32 and 1/64 element sizes is second order.

4.10 3D Hybrid 1x2x10 Duct: Specified Pressure Drop

In this section, a specified pressure drop in a simple 1x2x10 configuration is run with a variety of homogeneous blocks of the following topology: hexahedral, tetrahedral, wedge, and thexahedral. This analytical solution is given by an infinite series and is coded as the “1x2x10” user function. The simulation is run with an outer wall boundary condition with two open boundary conditions. The specified pressure drop is 0.016 over the 10 cm duct. The density and viscosity are $1.0\text{e-}3$ and $1.0\text{e-}4$, respectively. The simulation study is run a fixed Courant numbers with a mesh spacing ranging from 0.2 to 0.025. Figure Fig. 4.18 provides the standard velocity profile for the structured hexahedral and unstructured tetrahedral element type.




Fig. 4.18: Streamwise velocity profile for specified pressure drop flow; tetrahedral and hexahedral topology.

The simulation study employed a variety of elemental topologies of uniform mesh spacing as noted above. Figure Fig. 4.19 outlines the convergence in the L_2 norm using the low-order elemental CVFEM implementation using the recently changed tetrahedral and wedge element quadrature rules. Second-order accuracy is noted. Interestingly, the hexahedral and wedge topology provided nearly the same accuracy. Also, the tetrahedral accuracy was approximately four times greater. Finally, the Thexahedral topology proved to be second-order, however, provided very poor accuracy results.

4.11 3D Hybrid 1x1x1 Cube: Laplace

The standard Laplace operator is evaluated on the full set of low-order hybrid topologies (not including the pyramid). In this example, the temperature field is again,

$$T = \frac{\lambda}{4}(\cos(2a\pi x) + \cos(2a\pi y) + \cos(2a\pi z)). \quad (4.15)$$

Figure Fig. 4.20 represents the MMS field for temperature on a variety of mesh topologies. The thexahedral mesh is obtained from the standard uniform spacing tetrahedral mesh (not shown). The tetrahedral mesh shown is a tet-based conversion of the standard structured hexahedral mesh. This approach ensures that the number of nodes between the hexahedral and tetrahedral mesh are the same.

Figure Fig. 4.21 provides the L_2 norms, all of which are showing second-order accuracy. In Figure Fig. 4.22, the L_o error is shown. As indicated from the convergence plot, slight degradation in order-of-accuracy is noted for the thexahedral topology.




Fig. 4.19: L_2 error for the CVFEM scheme on a variety of element types.




Fig. 4.20: Temperature shadings for hexahedral, hexahedral, wedge, and tetrahedral topologies (clockwise from the upper left).

4.12 Fixed Wing Verification Problem

Introduction A basic verification of the actuator line theory can be done by considering a fixed, 2D wing in a uniform flow. Such a test can be done independently of the openFAST library and can be easily verified by a simple algebraic calculation.

```
actuator:
  type: ActLineSimple
  search_method: stk_kdtree
  search_target_part: Unspecified-2-HEX

  n_simpleblades: 1
  debug_output: false
  n_turbines_glob: 0
  Blade0:
    num_force_pts_blade: 20
    epsilon: [3.0, 3.0, 3.0]
    p1: [-25, -4, 0]
    p2: [-25, 4, 0]
    p1_zero_alpha_dir: [1, 0, 0]
    chord_table: [1.0]
    twist_table: [0.0]
    aoa_table: [-180, 0, 180]
    cl_table: [-19.739208802178716, 0, 19.739208802178716]
    cd_table: [0]
```

The fixed wing is defined between points \mathbf{p}_1 and \mathbf{p}_2 given the chord length and blade twist defined in **:code-block:‘chord_table’** and **:code-block:‘twist_table’**. The direction $\mathbf{p1}_{0\alpha}$ corresponding to the zero degree angle of




Fig. 4.21: L_2 norms for the full set of hybrid Laplace MMS study.




Fig. 4.22: L_2 norms for the full set of hybrid Laplace MMS study.




Fig. 4.23: Schematic of the fixed wing airfoil problem.

An example of the fixed wing specification in the Nalu-Wind input file is shown below. The actuator type can be `:code-block:'ActLineSimpleNGP'` for the NGP version and `:code-block:'ActLineSimple'` for the non-NGP version (soon to be deprecated).

attack is given in `:code-block:'p1_zero_alpha_dir'`. The lift coefficients C_L and drag coefficients C_D are tabulated in the `:code-block:'cl_table'` and `:code-block:'cd_table'` parameters, respectively, as functions of the angle of attack α in `:code-block:'aoa_table'`.

The lift L and drag D on the fixed wing can be calculated by infinite 2D airfoil theory using the formulas:

$$L = \frac{1}{2} \rho U^2 C_L(\alpha) S \quad (4.16)$$

$$D = \frac{1}{2} \rho U^2 C_D(\alpha) S \quad (4.17)$$

In equations (4.16) and (4.17), the area of the airfoil is given by S , the density is given by ρ , and the wind speed is given by U .

For the verification problem, a simple 2D, extruded blade is used with span $l = 8\text{m}$ and chord $c = 1\text{m}$, giving an area $S = 8\text{m}^2$. An isotropic force spreading value of $\varepsilon = 3$ is used, along with 20 blade stations. The wind conditions are given by $\rho = 1.0\text{ kg/m}^3$ and $U = 2\text{ m/s}$.

The aerodynamic properties of the fixed wing are given by the linear lift coefficient

$$C_L = 2\pi\alpha \quad (4.18)$$

and zero drag

$$C_D = 0. \quad (4.19)$$

A comparison of total lift force calculated Nalu-Wind against the 2D airfoil theory is shown in Figure `_fw_bladeresults`. As expected, the total lift force varies linearly with the angle of attack, and the agreement between theory and Nalu-Wind is good. Differences between the two methods were seen to be less than 0.1% over the range $0 \leq \alpha \leq 5$ degrees.

4.13 Actuator line simulations coupled to OpenFAST

We test the implementation of the actuator line algorithm in Nalu-Wind coupled to OpenFAST by performing a simulation of a flow past an elliptic wing at a constant angle of attack. We compare the solution from the coupled simulation to that using lifting line theory [KP02].

The elliptic wing is modeled using OpenFAST, a aero-hydro-servo-elastic tool to model wind turbines developed by the National Renewable Energy Laboratory (NREL). A static wind turbine model was created in OpenFAST with just one elliptic wing and all other systems including structural deformation, controls, etc. are turned off. The elliptic wing simulated in this work is an infinitesimally thin wing with a maximum chord (c_0) of 1m and an aspect ratio (b/c_0) of




Fig. 4.24: Comparison of the total lift force from Nalu-Wind and from 2D airfoil theory.

10.0. The lift-curve slope ($dC_l/d\alpha$) of all airfoil sections on the wing is assumed to be 2π with no pressure or viscous drag. Using lifting line theory [KP02], the loads on the elliptic wing are

Area $S =$

$$\pi \frac{c_0}{2} \frac{b}{2},$$

Maximum circulation $\Gamma_{\max} =$

$$\frac{2bU_\infty(\alpha - \alpha_{L0})}{1 + 4b/2\pi c_0},$$

Lift coefficient $C_L \equiv$

$$\frac{L}{0.5\rho U_\infty^2 S} = \frac{\pi}{2} \frac{b}{S} \frac{\Gamma_{\max}}{U_\infty},$$

Lift coefficient $C_D \equiv$

$$\frac{D}{0.5\rho U_\infty^2 S} = \frac{\pi}{4S} \frac{\Gamma_{\max}^2}{U_\infty^2},$$

Constant induced downwash $w_i =$

$$\frac{\Gamma_{\max}}{2b}.$$

$$\begin{aligned}\text{Span } b &= 10m \\ \text{Max chord } c_0 &= 1.0m \\ \text{Angle of attack } \alpha &= 7^\circ \\ U_\infty &= 10.0m/s \\ \text{Reynolds number based on chord} &= 0.66M \\ \text{Number of actuator points across span} &= 50\end{aligned}\tag{4.20}$$

The flow past the elliptic wing is simulated in a domain of size $4b \times 3b \times 3b$. Some parameters of the simulation are described in equation (4.20). As described in the section [Wind Turbine Modeling](#), the actuator line algorithm solves the momentum equation with a body force term spread to the nodes where ϵ is the spreading width. It is necessary to maintain a constant ϵ to observe convergence of the solution with grid refinement. However, we do expect the solution from the actuator line algorithm to be closer to that from lifting line theory with reducing ϵ . Hence, we perform five numerical simulations with grid resolutions as shown in table shown below. Simulations *a,b,c* use $\epsilon = 1m$ and *d,e* use $\epsilon = 0.5m$. We expect to see grid convergence with simulations *a,b,c* while we expect simulations *d,e* to predict a solution closer to the lifting line solution compared to simulations *a,b,c*.

Case	$\Delta x/c_0$	Δt	ϵ/Δ
<i>a</i>	0.125	0.00125	8.0
<i>b</i>	0.25	0.0025	4.0
<i>c</i>	0.5	0.005	2.0
<i>d</i>	0.125	0.00125	4.0
<i>e</i>	0.25	0.0025	2.0

The data shown in [Fig. 4.25](#) - [Fig. 4.29](#) are computed purely using output from OpenFAST. Unfortunately OpenFAST can only output data at a maximum of 9 stations along the blade. For this specific work, I had designed the aerodynamics module (AeroDyn) inside OpenFAST to use 18 stations to compute the forces along the blade. However, the mesh mapping algorithm in OpenFAST is used to interpolate the forces per unit length along the blade into discrete point forces at 50 actuator points along the blade as described in equation (4.20).

[Fig. 4.25](#)-[Fig. 4.26](#) shows the comparison of lift and drag coefficient predicted by the actuator line simulations to the solution from lifting line theory. Simulations *d* and *e* are closer to the lifting line solution compared to *a,b,c* because of the smaller ϵ . Simulations *a,b,c* show grid convergence since they use the same ϵ . [Fig. 4.27](#)-[Fig. 4.28](#) show similar results through the span wise distribution of the lift and drag per unit length along the blade. [Fig. 4.29](#) shows the comparison of the predicted angle of attack on the blade to the constant angle attack predicted by the lifting line theory. As expected, the agreement with the lifting line theory is much better near the mid-span region compared to the wing tips.

4.14 Open Boundary Condition With Outflow Thermal Stratification

In situations with significant thermal stratification at the outflow of the domain, the standard open boundary condition alone is not adequate because it requires the specification of motion pressure at the boundary, and this is not known *a priori*. Two solutions to this problem are: 1) to use the global mass flow rate correction option, or 2) to use the standard open boundary condition in which the buoyancy term uses a local time-averaged reference value, rather than a single reference value.

We test these open boundary condition options on a simplified stratified flow through a channel with slip walls. The flow entering the domain is non-turbulent and uniformly 8 m/s. The temperature linearly varies from 300 K to 310 K from the bottom to top of the channel with compatible, opposite-sign heat flux on the two walls to maintain this profile. The Boussinesq buoyancy option is used, and the density is set constant to 1.17804 kg/m³. This density is compatible with the reference pressure of 101325 Pa and a reference temperature of 300 K. The viscosity is set to 1.0e-5 Pa-s. *The flow should keep its inflow velocity and temperature profiles throughout the length of the domain.*



Fig. 4.25: Comparison of lift coefficient C_L for an elliptic wing simulated using actuator line algorithm to solution using lifting line theory.




Fig. 4.26: Comparison of drag coefficient C_D for an elliptic wing simulated using actuator line algorithm to solution using lifting line theory.




Fig. 4.27: Comparison of lift coefficient C_L for an elliptic wing simulated using actuator line algorithm to solution using lifting line theory. Results are only shown at 9 different stations along the blade that are output from OpenFAST.

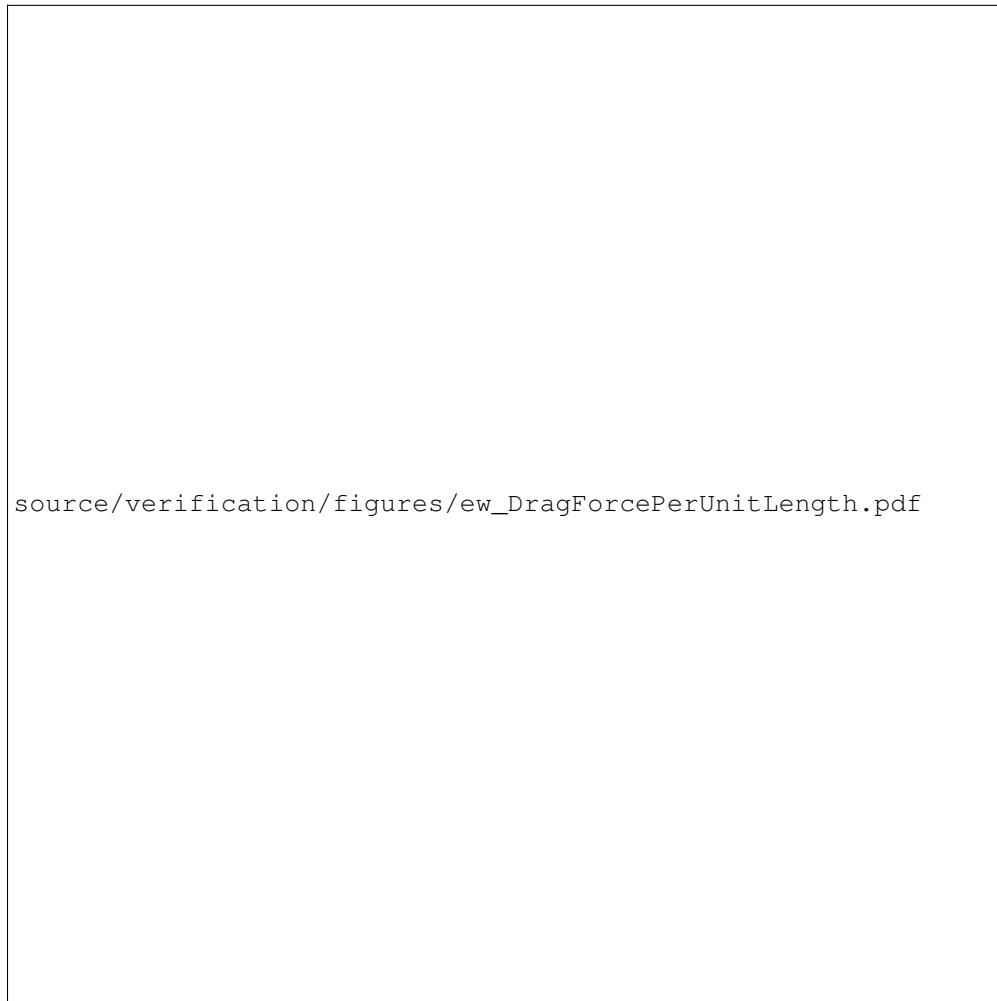


Fig. 4.28: Comparison of drag coefficient C_D for an elliptic wing simulated using actuator line algorithm to solution using lifting line theory. Results are only shown at 9 different stations along the blade that are output from OpenFAST.




Fig. 4.29: Comparison of angle of attack distribution on an elliptic wing simulated using actuator line algorithm to solution using lifting line theory. Results are only shown at 9 different stations along the blade that are output from OpenFAST.

The domain is 3000 m long, 1000 m tall, and 20 m wide with $300 \times 100 \times 2$ elements. The upper and lower boundaries are symmetry with the specified normal gradient of temperature option used such that the gradient matches the initial temperature profile with its gradient of 0.01 K/m. Flow enters from the left and exits on the right. The remaining boundaries are periodic.

We test the problem on three configurations: 1) using the standard open boundary condition, 2) using the global-mass-flow-rate-correction option, and 3) using the standard open boundary condition with a local moving-time-averaged reference temperature in the Boussinesq buoyancy term.

Figure Fig. 4.30 shows the across-channel profile of outflow streamwise velocity. It is clear that in configuration 1, the velocity is significantly distorted from the correct solution. Configurations 2 and 3 remedy the problem. However, if we reduce the range of the x-axis, as shown in Figure Fig. 4.31, we see that configuration 3, the use of the standard open boundary condition with a local moving-time-averaged Boussinesq reference temperature, provides a superior solution in this case. In Figure, Fig. 4.32, we also see that configuration 1 significantly distorts the temperature from the correct solution.




Fig. 4.30: Outflow velocity profiles for the thermally stratified slip-channel flow.

We also verify that the global mass-flow-rate correction of configuration 2 is correcting the outflow mass flow rate properly. The output from Nalu-Wind showing the correction is correct and is shown as follows:

```
Mass Balance Review:
Density accumulation: 0
Integrated inflow: -188486.0356751138
Integrated open: 188486.035672821
Total mass closure: -2.29277e-06
A mass correction of: -2.86596e-09 occurred on: 800 boundary integration points:
Post-corrected integrated open: 188486.0356751139
```




Fig. 4.31: Outflow velocity profiles for the thermally stratified slip-channel flow considering only the case with the global mass-flow-rate correction and the standard open boundary with the local moving-time-averaged Boussinesq reference value.




Fig. 4.32: Outflow temperature profiles for the thermally stratified slip-channel flow.

4.15 Specified Normal Temperature Gradient Boundary Condition

The motivation for adding the ability to specify the boundary-normal temperature gradient is atmospheric boundary layer simulation in which the upper portion of the domain often contains a stably stratified layer with a temperature gradient that extends all the way to the upper boundary. The desire is for the simulation to maintain that gradient throughout the simulation duration.

Our test case is a laminar infinite channel with slip walls. In this case, the flow velocity is zero so the problem is simply a heat conduction through fluid. The density is fixed as constant, and there are no source terms including buoyancy.

This problem has an the analytical solution for the temperature profile across the channel:

$$T(t, z) = T(t_0, z_0) + \frac{-g_H - g_0}{H} \kappa_{eff}(t - t_0) + g_0(z - z_0) + \frac{-g_H - g_0}{2H} (z - z_0)^2, \quad (4.21)$$

where t_0 is the initial time; z_0 is the height of the lower channel wall; H is the channel height; g_0 and g_H are the wall-normal gradients of temperature at the lower and upper walls, respectively; κ_{eff} is the effective thermal diffusivity; and z is the distance in the cross-channel direction. The sign of the temperature gradients assumes that boundary normal points inward from the boundary. For this solution to hold, the initial solution must be that of (4.21) with $t = t_0$.

For all test cases, we use a domain that is 10 m x 10 m in the periodic (infinite) directions, and 100 m in the cross-channel (z) direction. We specify a constant density of 1 kg/m³, zero velocity, no buoyancy source term, a viscosity of 1 Pa-s, and a laminar Prandtl number of 1. No turbulence model is used. The value of $T(t_0, z_0)$ is 300 K.

4.15.1 Simple Linear Temperature Profile: Equal and Opposite Specified Temperature Gradients

A simple verification test that is representative of a stable atmospheric capping inversion is to compute the simple thermal channel with equal and opposite specified temperature gradients on each wall. By setting $g_H = -g_0$ in Equation (4.21), we are left with

$$T(z) = T(z_0) + g_0(z - z_0). \quad (4.22)$$

In other words, if we set the initial temperature profile to that of (4.22), with $g_H = -g_0$, the profile should remain fixed for all time. In this case, we set $g_0 = 0.01$ K/m and $g_H = -0.01$ K/m.

We use a mesh that 2 elements wide in the periodic directions and 20 elements across the channel. We simulate a long time period of 25,000 s. Figure Fig. 4.33 shows that the computed and analytical solutions agree.

4.15.2 Parabolic Temperature Profile: Equal Specified Temperature Gradients

Next, we verify the specified normal temperature gradient boundary condition option by computing the simple thermal channel with equal specified temperature gradients, which yields the full time-dependent solution of Equation (4.21). Here, we set $g_0 = g_H = 0.01$ K/m.

We use meshes that are 2 elements wide in the periodic directions and 20, 40, and 80 elements across the channel. We simulate a long time period of 25,000 s. Figure Fig. 4.34 shows that the computed and analytical solutions agree. There is no apparent overall solution degradation on the coarser meshes.




Fig. 4.33: The analytical (black solid) and computed (red dashed) temperature profile from the case with $g_H = -g_0$ at $t = 25,000$ s.




Fig. 4.34: The analytical (black solid) and computed (colored) temperature profile from the case with $g_H = g_0$ at $t = 25,000$ s.

BIBLIOGRAPHY

- [McAl1991] McAlister, K. W. and R. K. Takahashi. *NACA 0015 Wing Pressure and Trailing Vortex Measurements*, NASA Technical Paper 3151. 1991.
- [Aft94] M. Aftosmis. Upwind method for simulation of viscous flow on adaptively refined meshes. *AIAA Journal*, 32(2):268–277, 1994.
- [Bau11] Mary C. Bautista. *Turbulence modelling of the atmospheric boundary layer over complex topography*. PhD thesis, École de Technologie Supérieure Université du Québec, Montreal, CA, October 2011.
- [BDM15] Mary C. Bautista, Louis Dufresne, and Christian Masson. Hybrid turbulence models for atmospheric flow; a proper comparison with rans models. *E3S Web of Conferences*, 2015. URL: https://www.e3s-conferences.org/articles/e3sconf/abs/2015/01/e3sconf_sowe2014_03001/e3sconf_sowe2014_03001.html.
- [CLM+12] M. J. Churchfield, S. Lee, P. K. Moriarty, L. A. Martinez, S. Leonardi, G. Vijayakumar, and J. G. Brasseur. A large-eddy simulation of wind-plant aerodynamics. In *Proceedings of the 50th AIAA Aerospace Sciences Meeting including the New Horizons Forum and Aerospace Exposition*. 9–12 Jan. 2012.
- [Dav97] L. Davidson. Large-eddy simulations: a note on the derivation of the equations for the subgrid turbulent kinetic energies. Technical Report, Chalmers University of Technology, Department of Thermo and Fluid Dynamics, 1997.
- [Dom06] S. Domino. Towards verification of formal time accuracy for a family of approximate projection methods using the method of manufactured solutions. In *Center for Turbulence Research Summer Proceedings*. 2006.
- [Dom08] S. Domino. A comparison of various equal-order interpolation methodologies using the method of manufactured solutions. In *Center for Turbulence Research Summer Proceedings*. 2008.
- [Dom10] S. Domino. Towards verification of sliding mesh algorithms for complex applications using mms. In *Center for Turbulence Research Summer Proceedings*. 2010.
- [Dom14] S. Domino. A comparison between low order and higher order low mach discretization approaches. In *Center for Turbulence Research Summer Proceedings*. 2014.
- [DNP98] F. Ducors, F. Nicoud, and T. Poinot. Wall-adapting local eddy-viscosity models for simulations in complex geometries. In *International Conference on Computational Conference*, volume 50. 1998.
- [Dye74] A. J. Dyer. A review of flux-profile relationships. *Boundary-Layer Meteorology*, 7:363–372, 1974.
- [EWS+10] H. Edwards, A. Williams, G. Sjaardema, D. Baur, and W. Cochran. Sierra toolkit computational mesh computational model. Technical Report SAND-20101192, Sandia National Laboratories, Albuquerque, NM, 2010.

- [G03] Tucker P. G. Differential equation-based wall distance computation for DES and RANS. *Journal of Computational Physics*, 190(1):229–248, September 2003. URL: [https://doi.org/10.1016/S0021-9991\(03\)00272-9](https://doi.org/10.1016/S0021-9991(03)00272-9).
- [HOM20] S. Haering, T. Oliver, and R. Moser. Active model split hybrid RANS/LES. *Journal of Fluid Mechanics*, 2020. Submitted. URL: <http://arxiv.org/abs/2006.13118>.
- [HOM19] Sigfried Haering, Todd Oliver, and Robert D. Moser. *Towards a Predictive Hybrid RANS/LES Framework*, chapter, pages. AIAA, 2019. URL: <https://arc.aiaa.org/doi/abs/10.2514/6.2019-0087>, arXiv:<https://arc.aiaa.org/doi/pdf/10.2514/6.2019-0087>, doi:10.2514/6.2019-0087.
- [HBH+03] M. Heroux, R. Bartlett, V. Howle, R. Hoekstra, J. Hu, T. Kolda, R. Lehoucq, K. Long, R. Pawlowski, E. Phipps, A. Salinger, J. Thornquist, R. Tuminaro, J. Willenbring, and A. Williams. An overview of trilinos. Technical Report SAND-20032927, Sandia National Laboratories, Albuquerque, NM, 2003.
- [Jas96] H. Jasek. Error analysis and estimation for the finite volume method with applications to fluid flow. In *Ph.D. Thesis, Imperial College*. 1996.
- [KV93] Y. Kallinderis and P. Vijayan. Adaptive refinement-coarsening scheme for three-dimensional unstructured meshes. *AIAA Journal*, 31(8):1440–1447, 1993.
- [KB89] Y. G. Kallinderis and J. R. Baron. Adaptive methods for a new navier-stokes algorithm. *AIAA Journal*, 27(1):37–43, 1989.
- [KP02] Joseph Katz and Allen Plotkin. *Low Speed Aerodynamics*. Cambridge University Press, second edition, 2002.
- [Mar05] M. Martinez. Comparison of galerkin and control volume finite element for advection-diffusion problems. *Int. J. Num. Meth. Fluids*, 50(3):347–376, 2005.
- [MartinezT17] Luis A Martínez-Tossas. *Large Eddy Simulations and Theoretical Analysis of Wind Turbine Aerodynamics Using an Actuator Line Model*. PhD thesis, Johns Hopkins University, Baltimore, MD USA, July 2017.
- [Mav00] D. J. Mavriplis. Adaptive meshing techniques for viscous flow calculations on mixed element unstructured meshes. *International Journal for Numerical Methods in Fluids*, 34(2):93–111, 2000.
- [MKL03] F. R. Menter, M. Kuntz, and R. Langtry. Ten years of industrial experience with the sst turbulence model. *Turb, Heat and Mass Trans*, 2003.
- [Moe84] C.-H. Moeng. A large-eddy-simulation model for the study of planetary boundary-layer turbulence. *J. Atmos. Sci.*, 41(13):2052–2062, 1984.
- [Pao82] S. Paolucci. On the filtering of sound waves from the navier-stokes equations. Technical Report SAND-828257, Sandia National Laboratories, Livermore, CA, December 1982.
- [RB78] R. G. Rehm and H. R. Baum. The equations of motion for thermally driven buoyant flows. *Journal of Research of the National Bureau of Standards*, 83:279, 1978.
- [RM84] R. Rogallo and P. Moin. Numerical simulation of turbulent flows. *Annual Review of Fluid Mechanics*, 16:99–137, 1984.
- [SR87] G. Schneider and M. Raw. Control volume finite element method for heat transfer and fluid flow using colocated variables - 1. computational procedure. *Numerical Heat Transfer*, 11(4):363–390, 1987.
- [SHZ91] F. Shakib, T. J. R. Hughes, and J. Zdenek. A new finite element formulation for computational fluids dynamics: the compressible euler and navier stokes equations. *Comp. Meth. in App. Mech and Engr.*, 89:141–219, 1991.
- [SorensenS02] Jens Nørkær Sørensen and Wen Zhong Shen. Numerical modeling of wind turbine wakes. *Journal of Fluids Engineering*, 124(2):393–399, 05 2002. URL: <http://dx.doi.org/10.1115/1.1471361>.

- [Tea16] SIERRA Thermal/Fluid Development Team. Sierra low mach module: fuego theory manual - version 4.42. Technical Report SAND2016-10163, Sandia National Laboratories, October 2016.
- [TDB05] S. Tieszen, S. Domino, and A. Black. Validation of a simple turbulence model suitable for closure of temporally-filtered navier-stokes equations using a helium plume. Technical Report SAND-20053210, Sandia National Laboratories, Albuquerque, NM, June 2005.

Symbols

-D, -debug
 nalux command line option, 15

-h, -help
 nalux command line option, 15

-i, -input-deck
 nalux command line option, 15

-o, -log-file
 nalux command line option, 15

-p, -pprint
 nalux command line option, 15

-v, -version
 nalux command line option, 15

A

abl_forcing
 Nalu-Wind Input File Parameter, 41

abl_forcing.momentum.computed
 Nalu-Wind Input File Parameter, 42

abl_forcing.momentum.heights
 Nalu-Wind Input File Parameter, 42

abl_forcing.momentum.relaxation_factor
 Nalu-Wind Input File Parameter, 42

abl_forcing.momentum.target_part_format
 Nalu-Wind Input File Parameter, 42

abl_forcing.momentum.velocity_x
 Nalu-Wind Input File Parameter, 42

abl_forcing.momentum.velocity_y
 Nalu-Wind Input File Parameter, 42

abl_forcing.momentum.velocity_z
 Nalu-Wind Input File Parameter, 42

abl_forcing.output_frequency
 Nalu-Wind Input File Parameter, 42

abl_forcing.search_expansion_factor
 Nalu-Wind Input File Parameter, 42

abl_forcing.search_method
 Nalu-Wind Input File Parameter, 42

abl_forcing.search_tolerance
 Nalu-Wind Input File Parameter, 42

abl_forcing.temperature.temperature
 Nalu-Wind Input File Parameter, 43

activate_aura
 Nalu-Wind Input File Parameter, 21

activate_memory_diagnostic
 Nalu-Wind Input File Parameter, 22

actuator
 Nalu-Wind Input File Parameter, 32

actuator.air_density
 Nalu-Wind Input File Parameter, 34

actuator.debug
 Nalu-Wind Input File Parameter, 33

actuator.dry_run
 Nalu-Wind Input File Parameter, 33

actuator.dt_fast
 Nalu-Wind Input File Parameter, 34

actuator.epsilon
 Nalu-Wind Input File Parameter, 34

actuator.epsilon_chord
 Nalu-Wind Input File Parameter, 34

actuator.epsilon_min
 Nalu-Wind Input File Parameter, 34

actuator.epsilon_tower
 Nalu-Wind Input File Parameter, 34

actuator.FAST_input_filename
 Nalu-Wind Input File Parameter, 34

actuator.n_every_checkpoint
 Nalu-Wind Input File Parameter, 34

actuator.n_turbines_glob
 Nalu-Wind Input File Parameter, 33

actuator.nacelle_area
 Nalu-Wind Input File Parameter, 34

actuator.nacelle_cd
 Nalu-Wind Input File Parameter, 34

actuator.num_force_pts_blade
 Nalu-Wind Input File Parameter, 34

actuator.num_force_pts_tower
 Nalu-Wind Input File Parameter, 34

actuator.num_swept_pts
 Nalu-Wind Input File Parameter, 35

actuator.restart_filename
 Nalu-Wind Input File Parameter, 34

actuator.search_method
 Nalu-Wind Input File Parameter, 33

actuator.simStart

Nalu-Wind Input File Parameter, 33
 actuator.t_end
 Nalu-Wind Input File Parameter, 33
 actuator.t_max
 Nalu-Wind Input File Parameter, 34
 actuator.t_start
 Nalu-Wind Input File Parameter, 33
 actuator.turb_id
 Nalu-Wind Input File Parameter, 35
 actuator.turbine_base_pos
 Nalu-Wind Input File Parameter, 34
 actuator.type
 Nalu-Wind Input File Parameter, 33
 automatic_decomposition_type
 Nalu-Wind Input File Parameter, 21

B

balance_node_iterations
 Nalu-Wind Input File Parameter, 22
 balance_node_target
 Nalu-Wind Input File Parameter, 22
 balance_nodes
 Nalu-Wind Input File Parameter, 22
 bc.target_name
 Nalu-Wind Input File Parameter, 24
 bc.wall_user_data
 Nalu-Wind Input File Parameter, 25
 boundary_conditions
 Nalu-Wind Input File Parameter, 24
 boundary_layer_statistics
 Nalu-Wind Input File Parameter, 43
 boundary_layer_statistics.compute_temperature
 Nalu-Wind Input File Parameter, 43
 boundary_layer_statistics.height_multiplier
 Nalu-Wind Input File Parameter, 44
 boundary_layer_statistics.minimum_height
 Nalu-Wind Input File Parameter, 43
 boundary_layer_statistics.output_frequency
 Nalu-Wind Input File Parameter, 43
 boundary_layer_statistics.process_u_tau_scaling
 Nalu-Wind Input File Parameter, 43
 boundary_layer_statistics.stats_output_file
 Nalu-Wind Input File Parameter, 43
 boundary_layer_statistics.target_name
 Nalu-Wind Input File Parameter, 43
 boundary_layer_statistics.time_filter_index
 Nalu-Wind Input File Parameter, 43
 boundary_layer_statistics.time_hist_output_frequency
 Nalu-Wind Input File Parameter, 43
 boundary_layer_statistics.wall_normal_direction
 Nalu-Wind Input File Parameter, 43

D

data_probes

Nalu-Wind Input File Parameter, 37
 data_probes.gzip_level
 Nalu-Wind Input File Parameter, 38
 data_probes.lidar_specifications
 Nalu-Wind Input File Parameter, 39
 data_probes.lidar_specifications.axis
 Nalu-Wind Input File Parameter, 39
 data_probes.lidar_specifications.beam_length
 Nalu-Wind Input File Parameter, 39
 data_probes.lidar_specifications.center
 Nalu-Wind Input File Parameter, 39
 data_probes.lidar_specifications.from_target_part
 Nalu-Wind Input File Parameter, 39
 data_probes.lidar_specifications.number_of_samples
 Nalu-Wind Input File Parameter, 39
 data_probes.lidar_specifications.points_along_line
 Nalu-Wind Input File Parameter, 39
 data_probes.lidar_specifications.scan_time
 Nalu-Wind Input File Parameter, 39
 data_probes.output_format
 Nalu-Wind Input File Parameter, 38
 data_probes.output_frequency
 Nalu-Wind Input File Parameter, 38
 data_probes.search_expansion_factor
 Nalu-Wind Input File Parameter, 38
 data_probes.search_method
 Nalu-Wind Input File Parameter, 38
 data_probes.search_tolerance
 Nalu-Wind Input File Parameter, 38
 data_probes.specifications
 Nalu-Wind Input File Parameter, 38
 data_probes.specifications.from_target_part
 Nalu-Wind Input File Parameter, 38
 data_probes.specifications.line_of_site_specifications
 Nalu-Wind Input File Parameter, 38
 data_probes.specifications.name
 Nalu-Wind Input File Parameter, 38
 data_probes.specifications.output_variables
 Nalu-Wind Input File Parameter, 39
 data_probes.specifications.plane_specifications
 Nalu-Wind Input File Parameter, 39
 data_probes.time_performance
 Nalu-Wind Input File Parameter, 38
 data_probes.write_coords
 Nalu-Wind Input File Parameter, 38
 data_probes.lidar_specifications.misc
 Nalu-Wind Input File Parameter, 39
 data_probes.lidar_specifications.target_courant
 Nalu-Wind Input File Parameter, 32
 data_probes.time_step_change_factor
 Nalu-Wind Input File Parameter, 32

E

equation_systems

Nalu-Wind Input File Parameter, 22
 equation_systems.max_iterations
 Nalu-Wind Input File Parameter, 22
 equation_systems.name
 Nalu-Wind Input File Parameter, 22
 equation_systems.solver_system_specification
 Nalu-Wind Input File Parameter, 22
 equation_systems.systems
 Nalu-Wind Input File Parameter, 23
 ExampleClass (C++ class), 82
 ExampleClass::~ExampleClass (C++ function), 82
 ExampleClass::AnotherMethod (C++ function), 83
 ExampleClass::DoNothing (C++ function), 82
 ExampleClass::DoSomething (C++ function), 82
 ExampleClass::ExampleClass (C++ function), 82
 ExampleClass::fAnswer (C++ member), 83
 ExampleClass::fQuestion (C++ member), 83
 ExampleClass::SomeProtectedMethod (C++ function), 83
 ExampleClass::VeryUsefulMethod (C++ function), 82

I

initial_conditions
 Nalu-Wind Input File Parameter, 23
 initial_conditions.constant
 Nalu-Wind Input File Parameter, 24
 initial_conditions.target_name
 Nalu-Wind Input File Parameter, 24
 initial_conditions.user_function
 Nalu-Wind Input File Parameter, 24

L

linear_solvers.bamg_coarsen_type
 Nalu-Wind Input File Parameter, 19
 linear_solvers.bamg_cycle_type
 Nalu-Wind Input File Parameter, 19
 linear_solvers.bamg_max_levels
 Nalu-Wind Input File Parameter, 19
 linear_solvers.bamg_num_sweeps
 Nalu-Wind Input File Parameter, 19
 linear_solvers.bamg_output_level
 Nalu-Wind Input File Parameter, 19
 linear_solvers.bamg_relax_order
 Nalu-Wind Input File Parameter, 19
 linear_solvers.bamg_relax_type
 Nalu-Wind Input File Parameter, 19
 linear_solvers.bamg_strong_threshold
 Nalu-Wind Input File Parameter, 19
 linear_solvers.kspace
 Nalu-Wind Input File Parameter, 18

linear_solvers.max_iterations
 Nalu-Wind Input File Parameter, 18
 linear_solvers.method
 Nalu-Wind Input File Parameter, 18
 linear_solvers.muelu_xml_file_name
 Nalu-Wind Input File Parameter, 19
 linear_solvers.name
 Nalu-Wind Input File Parameter, 18
 linear_solvers.output_level
 Nalu-Wind Input File Parameter, 19
 linear_solvers.preconditioner
 Nalu-Wind Input File Parameter, 18
 linear_solvers.recompute_preconditioner
 Nalu-Wind Input File Parameter, 19
 linear_solvers.reuse_preconditioner
 Nalu-Wind Input File Parameter, 19
 linear_solvers.summarize_muelu_timer
 Nalu-Wind Input File Parameter, 19
 linear_solvers.tolerance
 Nalu-Wind Input File Parameter, 18
 linear_solvers.type
 Nalu-Wind Input File Parameter, 18
 linear_solvers.write_matrix_files
 Nalu-Wind Input File Parameter, 19

M

material_properties
 Nalu-Wind Input File Parameter, 26
 material_properties.constant_specification
 Nalu-Wind Input File Parameter, 27
 material_properties.reference_quantities
 Nalu-Wind Input File Parameter, 27
 material_properties.specifications
 Nalu-Wind Input File Parameter, 27
 material_properties.specifications.name
 Nalu-Wind Input File Parameter, 27
 material_properties.specifications.type
 Nalu-Wind Input File Parameter, 27
 material_properties.target_name
 Nalu-Wind Input File Parameter, 27
 mesh
 Nalu-Wind Input File Parameter, 21
 mesh_motion
 Nalu-Wind Input File Parameter, 30
 mesh_motion.mesh_parts
 Nalu-Wind Input File Parameter, 30
 mesh_motion.motion
 Nalu-Wind Input File Parameter, 30
 mesh_motion.name
 Nalu-Wind Input File Parameter, 30
 mesh_transformation
 Nalu-Wind Input File Parameter, 29
 mesh_transformation.mesh_parts
 Nalu-Wind Input File Parameter, 29

mesh_transformation.motion
 Nalu-Wind Input File Parameter, 30
 mesh_transformation.name
 Nalu-Wind Input File Parameter, 29

N

Nalu-Wind Input File Parameter

abl_forcing, 41
 abl_forcing.momentum.computed, 42
 abl_forcing.momentum.heights, 42
 abl_forcing.momentum.relaxation_factor, 42
 abl_forcing.momentum.target_part_format, 42
 abl_forcing.momentum.velocity_x, 42
 abl_forcing.momentum.velocity_y, 42
 abl_forcing.momentum.velocity_z, 42
 abl_forcing.output_frequency, 42
 abl_forcing.search_expansion_factor, 42
 abl_forcing.search_method, 42
 abl_forcing.search_tolerance, 42
 abl_forcing.temperature.temperature, 43
 activate_aura, 21
 activate_memory_diagnostic, 22
 actuator, 32
 actuator.air_density, 34
 actuator.debug, 33
 actuator.dry_run, 33
 actuator.dt_fast, 34
 actuator.epsilon, 34
 actuator.epsilon_chord, 34
 actuator.epsilon_min, 34
 actuator.epsilon_tower, 34
 actuator.FAST_input_filename, 34
 actuator.n_every_checkpoint, 34
 actuator.n_turbines_glob, 33
 actuator.nacelle_area, 34
 actuator.nacelle_cd, 34
 actuator.num_force_pts_blade, 34
 actuator.num_force_pts_tower, 34
 actuator.num_swept_pts, 35
 actuator.restart_filename, 34
 actuator.search_method, 33
 actuator.simStart, 33
 actuator.t_end, 33
 actuator.t_max, 34
 actuator.t_start, 33
 actuator.turb_id, 35
 actuator.turbine_base_pos, 34
 actuator.type, 33
 automatic_decomposition_type, 21
 balance_node_iterations, 22
 balance_node_target, 22
 balance_nodes, 22
 bc.target_name, 24
 bc.wall_user_data, 25
 boundary_conditions, 24
 boundary_layer_statistics, 43
 boundary_layer_statistics.compute_temperature_s, 43
 boundary_layer_statistics.height_multiplier, 44
 boundary_layer_statistics.minimum_height, 43
 boundary_layer_statistics.output_frequency, 43
 boundary_layer_statistics.process_utau_statistics, 43
 boundary_layer_statistics.stats_output_file, 43
 boundary_layer_statistics.target_name, 43
 boundary_layer_statistics.time_filter_interval, 43
 boundary_layer_statistics.time_hist_output_frequency, 43
 boundary_layer_statistics.wall_normal_direction, 43
 data_probes, 37
 data_probes.gzip_level, 38
 data_probes.lidar_specifications, 39
 data_probes.lidar_specifications.axis, 39
 data_probes.lidar_specifications.beam_length, 39
 data_probes.lidar_specifications.center, 39
 data_probes.lidar_specifications.from_target_part, 39
 data_probes.lidar_specifications.number_of_samples, 39
 data_probes.lidar_specifications.points_along_line, 39
 data_probes.lidar_specifications.scan_time, 39
 data_probes.output_format, 38
 data_probes.output_frequency, 38
 data_probes.search_expansion_factor, 38
 data_probes.search_method, 38
 data_probes.search_tolerance, 38
 data_probes.specifications, 38
 data_probes.specifications.from_target_part, 38
 data_probes.specifications.line_of_site_specification, 38

data_probes.specifications.name, 38
data_probes.specifications.output_variables, 27
39
data_probes.specifications.plane_specifications, 39
data_probes.time_performance, 38
data_probes.write_coords, 38
dataprobess.lidar_specifications.misc, 39
dtctrl.target_courant, 32
dtctrl.time_step_change_factor, 32
equation_systems, 22
equation_systems.max_iterations, 22
equation_systems.name, 22
equation_systems.solver_system_specifications, 22
equation_systems.systems, 23
initial_conditions, 23
initial_conditions.constant, 24
initial_conditions.target_name, 24
initial_conditions.user_function, 24
linear_solvers.bamg_coarsen_type, 19
linear_solvers.bamg_cycle_type, 19
linear_solvers.bamg_max_levels, 19
linear_solvers.bamg_num_sweeps, 19
linear_solvers.bamg_output_level, 19
linear_solvers.bamg_relax_order, 19
linear_solvers.bamg_relax_type, 19
linear_solvers.bamg_strong_threshold, 19
linear_solvers.kspace, 18
linear_solvers.max_iterations, 18
linear_solvers.method, 18
linear_solvers.muelu_xml_file_name, 19
linear_solvers.name, 18
linear_solvers.output_level, 19
linear_solvers.preconditioner, 18
linear_solvers.recompute_preconditioner, 19
linear_solvers.reuse_preconditioner, 19
linear_solvers.summarize_muelu_timer, 19
linear_solvers.tolerance, 18
linear_solvers.type, 18
linear_solvers.write_matrix_files, 19
material_properties, 26
material_properties.constant_specifications, 27
material_properties.reference_quantities, 27
material_properties.specifications, 27
material_properties.specifications.name, 27
material_properties.specifications.type, 27
material_properties.target_name, 27
mesh, 21
mesh_motion, 30
mesh_motion.mesh_parts, 30
mesh_motion.motion, 30
mesh_motion.name, 30
mesh_transformation, 29
mesh_transformation.mesh_parts, 29
mesh_transformation.motion, 30
mesh_transformation.name, 29
name, 21
output, 30
output.compression_level, 31
output.output_data_base_name, 30
output.output_forced_wall_time, 31
output.output_frequency, 31
output.output_node_set, 31
output.output_start, 31
output.output_variables, 31
periodic_user_data, 26
polynomial_order, 22
post_processing, 40
post_processing.frequency, 40
post_processing.output_file_name, 40
post_processing.parameters, 40
post_processing.physics, 40
post_processing.target_name, 41
post_processing.type, 40
rebalance_mesh, 22
restart, 31
restart.compression_level, 31
restart.max_data_base_step_size, 31
restart.restart_data_base_name, 31
restart.restart_forced_wall_time, 31
restart.restart_frequency, 31
restart.restart_node_set, 31
restart.restart_start, 31
restart.restart_time, 31
search_target_part, 33
simulations, 44
solution_options, 29
solution_options.name, 29
solution_options.options, 29
solution_options.turbulence_model, 29
solve_frequency, 22
support_inconsistent_multi_state_restart, 22

time_int.name, 20	-h, -help, 15
time_int.realms, 20	-i, -input-deck, 15
time_int.second_order_accuracy, 20	-o, -log-file, 15
time_int.start_time, 20	-p, -pprint, 15
time_int.termination_step_count, 20	-v, -version, 15
time_int.termination_time, 20	name
time_int.time_step, 20	Nalu-Wind Input File Parameter, 21
time_int.time_step_count, 20	O
time_int.time_stepping_type, 20	output
Time_Integrators, 19	Nalu-Wind Input File Parameter, 30
time_step_control, 32	output.compression_level
transfers, 44	Nalu-Wind Input File Parameter, 31
turbulence_averaging, 35	output.output_data_base_name
turbulence_averaging.averaging_type, 35	Nalu-Wind Input File Parameter, 30
turbulence_averaging.forced_reset, 35	output.output_forced_wall_time
turbulence_averaging.specifications, 36	Nalu-Wind Input File Parameter, 31
turbulence_averaging.specifications.compute_favre_stress, 36	output.output_frequency
turbulence_averaging.specifications.compute_favre_stress, 36	Nalu-Wind Input File Parameter, 31
turbulence_averaging.specifications.compute_favre_tke, 36	output.output_start
turbulence_averaging.specifications.compute_lambda_cir, 37	Nalu-Wind Input File Parameter, 31
turbulence_averaging.specifications.compute_q_criterion, 36	output.output_variables
turbulence_averaging.specifications.compute_reynolds_stress, 36	Nalu-Wind Input File Parameter, 31
turbulence_averaging.specifications.compute_reynolds_stress, 36	P
turbulence_averaging.specifications.compute_sfs_stress, 36	periodic_user_data
turbulence_averaging.specifications.compute_temperature_resolved_flux, 36	Nalu-Wind Input File Parameter, 26
turbulence_averaging.specifications.compute_temperature_resolved_flux, 36	polynomial_order
turbulence_averaging.specifications.compute_temperature_sfs_flux, 36	Nalu-Wind Input File Parameter, 22
turbulence_averaging.specifications.compute_temperature_sfs_flux, 36	post_processing
turbulence_averaging.specifications.compute_tke, 36	Nalu-Wind Input File Parameter, 40
turbulence_averaging.specifications.compute_vorticity, 37	post_processing.frequency
turbulence_averaging.specifications.favre_average_variables, 36	Nalu-Wind Input File Parameter, 40
turbulence_averaging.specifications.name, 36	post_processing.output_file_name
turbulence_averaging.specifications.name, 36	Nalu-Wind Input File Parameter, 40
turbulence_averaging.specifications.reynolds_average_variables, 36	post_processing.parameters
turbulence_averaging.specifications.target_name, 36	Nalu-Wind Input File Parameter, 40
turbulence_averaging.time_filter_interval, 36	post_processing.physics
use_edges, 21	Nalu-Wind Input File Parameter, 40
naluX command line option	post_processing.target_name
-D, -debug, 15	Nalu-Wind Input File Parameter, 41
	post_processing.type
	Nalu-Wind Input File Parameter, 40
	R
	rebalance_mesh
	Nalu-Wind Input File Parameter, 22
	restart
	Nalu-Wind Input File Parameter, 31
	restart.compression_level
	Nalu-Wind Input File Parameter, 31
	restart.max_data_base_step_size

Nalu-Wind Input File Parameter, 31
 restart.restart_data_base_name
 Nalu-Wind Input File Parameter, 31
 restart.restart_forced_wall_time
 Nalu-Wind Input File Parameter, 31
 restart.restart_frequency
 Nalu-Wind Input File Parameter, 31
 restart.restart_node_set
 Nalu-Wind Input File Parameter, 31
 restart.restart_start
 Nalu-Wind Input File Parameter, 31
 restart.restart_time
 Nalu-Wind Input File Parameter, 31

S

search_target_part
 Nalu-Wind Input File Parameter, 33
 sierra::nalu::Actuator (C++ class), 77
 sierra::nalu::ActuatorLineFAST (C++ class), 77
 sierra::nalu::AuxFunction (C++ class), 77
 sierra::nalu::BoundaryLayerPerturbationAuxFunction (C++ class), 78
 sierra::nalu::ContinuityEquationSystem (C++ class), 67
 sierra::nalu::ConvectingTaylorVortexPressureAuxFunction (C++ class), 78
 sierra::nalu::ConvectingTaylorVortexPressureGradientAuxFunction (C++ class), 78
 sierra::nalu::ConvectingTaylorVortexVelocityAuxFunction (C++ class), 78
 sierra::nalu::DataProbePostProcessing (C++ class), 79
 sierra::nalu::EnthalpyEquationSystem (C++ class), 64
 sierra::nalu::EnthalpyEquationSystem::post_iter_work (C++ function), 65
 sierra::nalu::EnthalpyEquationSystem::solve_and_update (C++ function), 65
 sierra::nalu::EquationSystem (C++ class), 62
 sierra::nalu::EquationSystem::linsysWriteCounter (C++ member), 64
 sierra::nalu::EquationSystem::ngp_create_solver_function (C++ function), 63
 sierra::nalu::EquationSystem::ngpPecletFunctions (C++ member), 64
 sierra::nalu::EquationSystem::post_iter_work (C++ function), 63
 sierra::nalu::EquationSystem::post_iter_work_dep (C++ function), 63
 sierra::nalu::EquationSystem::postIterAlgDriver (C++ member), 64

sierra::nalu::EquationSystem::pre_iter_work (C++ function), 63
 sierra::nalu::EquationSystem::preIterAlgDriver_ (C++ member), 64
 sierra::nalu::EquationSystem::solution_update (C++ function), 63
 sierra::nalu::EquationSystem::solve_and_update (C++ function), 62
 sierra::nalu::EquationSystems (C++ class), 68
 sierra::nalu::EquationSystems::decoupledOversetGlobal (C++ member), 69
 sierra::nalu::EquationSystems::numOversetItersDefault (C++ member), 69
 sierra::nalu::EquationSystems::post_iter_work (C++ function), 68
 sierra::nalu::EquationSystems::postIterAlgDriver_ (C++ member), 69
 sierra::nalu::EquationSystems::pre_iter_work (C++ function), 68
 sierra::nalu::EquationSystems::preIterAlgDriver_ (C++ member), 69
 sierra::nalu::EquationSystems::solve_and_update (C++ function), 68
 sierra::nalu::HeatCondEquationSystem (C++ class), 66
 sierra::nalu::HeatCondEquationSystem::solve_and_update (C++ function), 66
 sierra::nalu::Hex27SCS (C++ class), 76
 sierra::nalu::Hex27SCV (C++ class), 76
 sierra::nalu::Hex8FEM (C++ class), 76
 sierra::nalu::HexSCS (C++ class), 76
 sierra::nalu::HexSCV (C++ class), 76
 sierra::nalu::HypreDirectSolver (C++ class), 74
 sierra::nalu::HypreDirectSolver::destroyLinearSolver (C++ function), 74
 sierra::nalu::HypreDirectSolver::getType (C++ function), 74
 sierra::nalu::HypreDirectSolver::parMat_ (C++ member), 75
 sierra::nalu::HypreDirectSolver::parRhs_ (C++ member), 75
 sierra::nalu::HypreDirectSolver::parSln_ (C++ member), 75
 sierra::nalu::HypreDirectSolver::set_initialize_solver (C++ function), 74
 sierra::nalu::HypreDirectSolver::solve (C++ function), 74
 sierra::nalu::HypreLinearSolverConfig (C++ class), 76
 sierra::nalu::HypreLinearSolverConfig::load (C++ function), 76
 sierra::nalu::HypreLinearSystem (C++

class), 70
 sierra::nalu::HypreLinearSystem::applyDirichletBcalsu::KovasznyVelocityAuxFunction
 (C++ function), 71
 sierra::nalu::HypreLinearSystem::buildDirichletNodeGraphLinearRampMeshDisplacementAuxFunction
 (C++ function), 70
 sierra::nalu::HypreLinearSystem::copy_hyperto::LinearSolver (C++ class), 61,
 (C++ function), 71
 sierra::nalu::HypreLinearSystem::HypreLinearSystem::activeMueLu
 (C++ function), 70
 sierra::nalu::HypreLinearSystem::HypreLinearSystem::destroyLinearSolver
 (C++ class), 72
 sierra::nalu::HypreLinearSystem::HypreLinearSystem::skipRows_precond
 (C++ member), 72
 sierra::nalu::HypreLinearSystem::HypreLinearSystem::solvePrinterConfig
 (C++ member), 72
 sierra::nalu::HypreLinearSystem::HypreLinearSystem::getType
 (C++ member), 72
 sierra::nalu::HypreLinearSystem::HypreLinearSystem::name_ (C++
 (C++ member), 72
 sierra::nalu::HypreLinearSystem::HypreLinearSystem::recomputePreconditioner
 (C++ member), 72
 sierra::nalu::HypreLinearSystem::HypreLinearSystem::HypreGlobalPreconditioner
 (C++ member), 72
 sierra::nalu::HypreLinearSystem::HypreLinearSystem::zero_timer_precond
 (C++ member), 72
 sierra::nalu::HypreLinearSystem::HypreLinearSystem::SolverConfig_ (C++ owned_
 (C++ member), 72
 sierra::nalu::HypreLinearSystem::HypreLinearSystem::rowsCofied::reuseLinSysIfPoss
 (C++ member), 72
 sierra::nalu::HypreLinearSystem::HypreLinearSystem::rows (C++ class), 75
 (C++ member), 72
 sierra::nalu::HypreLinearSystem::HypreLinearSystem::sumDof_
 (C++ member), 72
 sierra::nalu::HypreLinearSystem::HypreLinearSystem::overRowsMap_
 (C++ member), 72
 sierra::nalu::HypreLinearSystem::HypreLinearSystem::periodic_node_to_hypre_id_
 (C++ member), 72
 sierra::nalu::HypreLinearSystem::HypreLinearSystem::skippedRowsMap_
 (C++ member), 72
 sierra::nalu::HypreLinearSystem::loadComplete (C++ member), 75
 (C++ function), 70
 sierra::nalu::HypreLinearSystem::resetRows (C++ member), 75
 (C++ function), 71
 sierra::nalu::HypreLinearSystem::solve (C++ function), 71
 sierra::nalu::HypreLinearSystem::sumInto (C++ function), 71
 sierra::nalu::HypreLinearSystem::zeroSystem (C++ function), 71
 (C++ function), 71
 sierra::nalu::InputOutputRealm (C++ class), 60
 sierra::nalu::KovasznyPressureAuxFunction (C++ class), 78
 sierra::nalu::KovasznyPressureGradientAuxFunction (C++ function), 64

```

sierra::nalu::MassFractionEquationSystem (C++ class), 66
sierra::nalu::MassFractionEquationSystem::solve_and_update (C++ function), 66
sierra::nalu::MasterElement (C++ class), 76
sierra::nalu::MixtureFractionEquationSystem (C++ class), 67
sierra::nalu::MixtureFractionEquationSystem::solve_and_update (C++ function), 67
sierra::nalu::MomentumEquationSystem (C++ class), 67
sierra::nalu::ProjectedNodalGradientEquationSystem (C++ class), 67
sierra::nalu::ProjectedNodalGradientEquationSystem::solve_and_update (C++ function), 67
sierra::nalu::PyrSCS (C++ class), 76
sierra::nalu::PyrSCV (C++ class), 76
sierra::nalu::Quad3DSCS (C++ class), 76
sierra::nalu::Quad42DSCS (C++ class), 77
sierra::nalu::Quad42DSCV (C++ class), 77
sierra::nalu::Quad93DSCS (C++ class), 76
sierra::nalu::Realm (C++ class), 59
sierra::nalu::Realm::bcPartVec_ (C++ member), 60
sierra::nalu::Realm::check_job (C++ function), 59
sierra::nalu::Realm::hypreGlobalId_ (C++ member), 60
sierra::nalu::Realm::hypreILower_ (C++ member), 60
sierra::nalu::Realm::hypreIsActive_ (C++ member), 60
sierra::nalu::Realm::hypreIUpper_ (C++ member), 60
sierra::nalu::Realm::hypreNumNodes_ (C++ member), 60
sierra::nalu::Realm::set_hypre_global_ids (C++ function), 59
sierra::nalu::Realms (C++ class), 60
sierra::nalu::ShearStressTransportEquationSystem (C++ class), 65
sierra::nalu::ShearStressTransportEquationSystem::solve_and_update (C++ function), 66
sierra::nalu::ShearStressTransportEquationSystem::solve_and_update (C++ function), 65
sierra::nalu::Simulation (C++ class), 59
sierra::nalu::SinMeshDisplacementAuxFunction (C++ class), 79
sierra::nalu::SolutionNormPostProcessingsierra::nalu::SteadyTaylorVortexMomentumSrcElemSupp (C++ class), 78
sierra::nalu::SteadyTaylorVortexMomentumSrcNodeSupp (C++ class), 78
sierra::nalu::SteadyTaylorVortexPressureAuxFunction (C++ class), 78
sierra::nalu::SteadyTaylorVortexVelocityAuxFunction (C++ class), 78
sierra::nalu::SteadyThermal3dContactAuxFunction (C++ class), 78
sierra::nalu::SteadyThermal3dContactDtDxAuxFunction (C++ class), 78
sierra::nalu::SteadyThermal3dContactSrcElemKernel (C++ class), 78
sierra::nalu::SteadyThermal3dContactSrcElemKernel::solve_and_update (C++ function), 79
sierra::nalu::SteadyThermal3dContactSrcElemSuppAlg (C++ class), 79
sierra::nalu::SteadyThermalContact3DSrcNodeSuppAlg (C++ class), 79
sierra::nalu::SteadyThermalContactAuxFunction (C++ class), 79
sierra::nalu::SteadyThermalContactSrcElemSuppAlg (C++ class), 79
sierra::nalu::SteadyThermalContactSrcNodeSuppAlg (C++ class), 79
sierra::nalu::SurfaceForceAndMomentAlgorithm (C++ class), 79
sierra::nalu::SurfaceForceAndMomentWallFunctionAlgorithm (C++ class), 80
sierra::nalu::TetSCS (C++ class), 76
sierra::nalu::TetSCV (C++ class), 76
sierra::nalu::TimeIntegrator (C++ class), 60
sierra::nalu::TpetraLinearSolver (C++ class), 73
sierra::nalu::TpetraLinearSolver::destroyLinearSolver (C++ function), 73
sierra::nalu::TpetraLinearSolver::getType (C++ function), 74
sierra::nalu::TpetraLinearSolver::residual_norm (C++ function), 73
sierra::nalu::TpetraLinearSolver::setMueLuWork (C++ function), 73
sierra::nalu::TpetraLinearSolver::setMueLuWorkAndUpdate (C++ function), 73
sierra::nalu::TpetraLinearSolver::solve (C++ function), 74
sierra::nalu::TpetraLinearSolver::TpetraLinearSolver (C++ function), 73
sierra::nalu::TpetraLinearSolverConfig (C++ class), 76
sierra::nalu::TpetraLinearSystem (C++ class), 61, 69
sierra::nalu::TpetraLinearSystem::resetRows (C++ function), 62, 70

```

[sierra::nalu::Transfer \(C++ class\), 62](#)
[sierra::nalu::Transfers \(C++ class\), 62](#)
[sierra::nalu::Tri32DSCS \(C++ class\), 77](#)
[sierra::nalu::Tri32DSCV \(C++ class\), 77](#)
[sierra::nalu::Tri3DSCS \(C++ class\), 76](#)
[sierra::nalu::TurbKineticEnergyEquationSystem \(C++ class\), 65](#)
[sierra::nalu::TurbKineticEnergyEquationSystem::solve_and_update \(C++ function\), 65](#)
[sierra::nalu::TurbulenceAveragingPostProtocols \(C++ class\), 79](#)
[sierra::nalu::TurbulenceAveragingPostProtocols::TurbulenceAveragingType \(C++ enum\), 79](#)
[sierra::nalu::TurbulenceAveragingPostProtocols::TurbulenceAveragingType::EXPONENTIAL \(C++ enumerator\), 79](#)
[sierra::nalu::TurbulenceAveragingPostProtocols::TurbulenceAveragingType::MOVING_AVERAGE \(C++ enumerator\), 79](#)
[sierra::nalu::WedSCS \(C++ class\), 76](#)
[sierra::nalu::WedSCV \(C++ class\), 76](#)
[simulations](#)
 Nalu-Wind Input File Parameter, 44
[solution_options](#)
 Nalu-Wind Input File Parameter, 29
[solution_options.name](#)
 Nalu-Wind Input File Parameter, 29
[solution_options.options](#)
 Nalu-Wind Input File Parameter, 29
[solution_options.turbulence_model](#)
 Nalu-Wind Input File Parameter, 29
[solve_frequency](#)
 Nalu-Wind Input File Parameter, 22
[support_inconsistent_multi_state_restart](#)
 Nalu-Wind Input File Parameter, 22

T

[time_int.name](#)
 Nalu-Wind Input File Parameter, 20
[time_int.realms](#)
 Nalu-Wind Input File Parameter, 20
[time_int.second_order_accuracy](#)
 Nalu-Wind Input File Parameter, 20
[time_int.start_time](#)
 Nalu-Wind Input File Parameter, 20
[time_int.termination_step_count](#)
 Nalu-Wind Input File Parameter, 20
[time_int.termination_time](#)
 Nalu-Wind Input File Parameter, 20
[time_int.time_step](#)
 Nalu-Wind Input File Parameter, 20
[time_int.time_step_count](#)
 Nalu-Wind Input File Parameter, 20
[time_int.time_stepping_type](#)
 Nalu-Wind Input File Parameter, 20
[Time_Integrators](#)

[Nalu-Wind Input File Parameter, 19](#)
[time_step_control](#)
 Nalu-Wind Input File Parameter, 32
[transfers](#)
 Nalu-Wind Input File Parameter, 44
[turbulence_averaging](#)
 Nalu-Wind Input File Parameter, 35
[turbulence_averaging_type](#)
 Nalu-Wind Input File Parameter, 35
[turbulence_averaging.forced_reset](#)
 Nalu-Wind Input File Parameter, 35
[turbulence_averaging.type](#)
 Nalu-Wind Input File Parameter, 36
[turbulence_averaging.type.specifications](#)
 Nalu-Wind Input File Parameter, 36
[turbulence_averaging.type.specifications.compute_favre_tke](#)
 Nalu-Wind Input File Parameter, 36
[turbulence_averaging.type.specifications.compute_favre_tke_lambda](#)
 Nalu-Wind Input File Parameter, 37
[turbulence_averaging.type.specifications.compute_q_crit](#)
 Nalu-Wind Input File Parameter, 36
[turbulence_averaging.type.specifications.compute_resolve](#)
 Nalu-Wind Input File Parameter, 36
[turbulence_averaging.type.specifications.compute_reynolds](#)
 Nalu-Wind Input File Parameter, 36
[turbulence_averaging.type.specifications.compute_sfs_str](#)
 Nalu-Wind Input File Parameter, 36
[turbulence_averaging.type.specifications.compute_tempera](#)
 Nalu-Wind Input File Parameter, 36
[turbulence_averaging.type.specifications.compute_tempera](#)
 Nalu-Wind Input File Parameter, 36
[turbulence_averaging.type.specifications.compute_tke](#)
 Nalu-Wind Input File Parameter, 36
[turbulence_averaging.type.specifications.compute_vortic](#)
 Nalu-Wind Input File Parameter, 37
[turbulence_averaging.type.favre_average_v](#)
 Nalu-Wind Input File Parameter, 36
[turbulence_averaging.type.name](#)
 Nalu-Wind Input File Parameter, 36
[turbulence_averaging.type.reynolds_averag](#)
 Nalu-Wind Input File Parameter, 36
[turbulence_averaging.type.target_name](#)
 Nalu-Wind Input File Parameter, 36
[turbulence_averaging.time_filter_interval](#)
 Nalu-Wind Input File Parameter, 36

U

[use_edges](#)
 Nalu-Wind Input File Parameter, 21