

OpenSBLI: A framework for the automated derivation and parallel execution of finite difference solvers on a range of computer architectures

Christian T. Jacobs^{a,*}, Satya P. Jammy^a, Neil D. Sandham^a

^a*Aerodynamics and Flight Mechanics Group, Faculty of Engineering and the Environment, University of Southampton, University Road, Southampton, SO17 1BJ, United Kingdom*

Abstract

Exascale computing will feature novel and potentially disruptive hardware architectures. Exploiting these to their full potential is non-trivial. Numerical modelling frameworks involving finite difference methods are currently limited by the ‘static’ nature of the hand-coded discretisation schemes and repeatedly may have to be re-written to run efficiently on new hardware. In contrast, OpenSBLI uses code generation to derive the model’s code from a high-level specification. Users focus on the equations to solve, whilst not concerning themselves with the detailed implementation. Source-to-source translation is used to tailor the code and enable its execution on a variety of hardware.

Keywords: High-Performance Computing, Code Generation, Computational Fluid Dynamics, Finite Difference Methods, Graphics Processing Units

1. Introduction

High Performance Computing (HPC) systems and architectures are evolving rapidly. Traditional single processor-based CPU clusters are moving towards multi-core/multi-threaded CPUs. At the same time new architectures based on many-core processors such as graphics processing units (GPUs) and

*Corresponding author.
E-mail address: C.T.Jacobs@soton.ac.uk

6 Intel’s Xeon Phi are emerging as important systems and further developments
7 are expected with energy-efficient designs from ARM and IBM. According to
8 the IT industry, such advances are expected to deliver compute hardware ca-
9 pable of exascale-performance (i.e. 10^{18} floating-point operations per second)
10 by 2018 [1]. Yet many frameworks aimed at computational/numerical mod-
11 elling are currently not ready to exploit such new and potentially disruptive
12 technologies.

13 Traditional approaches to numerical model development involve the pro-
14 duction of static, hand-written code to perform the numerical discretisation
15 and solution of the governing equations. Normally this is written in a lan-
16 guage such as C or Fortran that is considerably less abstract when compared
17 to a near-mathematical domain specific language. Explicitly inserting the
18 necessary calls to MPI or OpenMP libraries enables the execution of the
19 code on multi-core or multi-thread hardware. However, should a user wish
20 to run the code on alternative platforms such as GPUs, they would likely
21 need to re-write large sections of the code, including calls to new libraries
22 such as CUDA or OpenCL, and optimise it for that particular hardware
23 backend [2]. As HPC hardware evolves, an increasing burdon faced by com-
24 putational scientists becomes apparent; in order to keep up with trends in
25 HPC, not only must a model developer be a domain specialist in their area
26 of study, but also an expert in numerical algorithms, software engineering,
27 and parallel computing paradigms [3, 4].

28 One way to address this issue is to introduce a separation of concerns
29 using high level abstractions, such as domain specific languages (DSLs) and
30 active libraries [4, 5, 6, 7, 8]. This paradigm shift allows a domain specialist to
31 describe their problem as a high-level, near-mathematical specification. The
32 task of taking this specification and transforming it into executable computer
33 code can then be handled in the subsequent abstraction layer; unlike the
34 traditional approach of hand-writing the C/Fortran code that discretises the
35 governing equations, this layer generates the code automatically from the
36 problem specification. Finally, the generated code can be readily targetted
37 towards a specific hardware platform through source-to-source translation.
38 Hence, domain specialists focus on the equations they wish to solve and the
39 setup of their problem, whilst the parallel computing experts can introduce
40 support for new backends as they become available. At no point does the
41 code have to undergo a fundamental re-write if the desired backend changes.
42 Use of such strategies can have significant benefits for the productivity of
43 both the user and developer, by removing the need to spend time re-writing

code and/or the problem specification [5].

Given the motivation for the use of automated solution techniques, in this paper we present a new framework, OpenSBLI, for the automated derivation and parallel execution of finite difference-based models. This is an open-source release of the recent developments in the SBLI codebase developed at the University of Southampton, involving the replacement of SBLI’s Fortran-based core with flexible Python-based code generation capabilities, and the coupling of SBLI to the OPS active library [9, 10, 11, 12] which targets the generated code towards a particular backend using source-to-source translation. Currently, OpenSBLI can generate OPS-compliant C code to discretise and solve the governing equations, using arbitrary-order central finite difference schemes and a choice of either the forward Euler scheme or a third-order Runge-Kutta time-stepping scheme. OpenSBLI then uses OPS to produce code targetted towards different backends. It is worth noting that backend APIs such as OpenMP (version 4.0 and above) are also capable of running on CPU, GPU and Intel Xeon Phi architectures, for example. However, currently OPS has no support for OpenMP version 4.0 and above. Moreover, codes that are written by hand in OpenMP would still potentially need to be re-written if different algorithms or equations were to be considered. Thus, the benefits of code generation still play a crucial role here, regardless of which backend is chosen.

The application of SBLI has so-far concentrated on problems in aeronautics and aeroacoustics, in particular looking at shock-boundary layer interactions (see e.g. [13, 14, 15, 16] and the references therein for more details). While such applications entail solving the 3D compressible Navier-Stokes equations, in principle other equations expressible in Einstein notation and solved using finite differences are also supported by the new code generation functionality, highlighting another advantage of such a flexible approach to numerical model development. Note also that while OpenSBLI does not yet feature shock-capturing schemes and Large Eddy Simulation models (unlike the legacy SBLI code), these will be implemented in the future as part of the project’s roadmap. The main purposes of this initial release is the algorithmic changes to legacy SBLI’s core.

Details the abstraction and design principles employed by OpenSBLI are given in Section 2. Section 3 details three verification and validation test cases that were used to check the correctness of the implementation. The paper finishes with some concluding remarks in Section 4.

2. Design

Legacy versions of SBLI comprise static hand-written Fortran code, parallelised with MPI, that implements a fourth-order central differencing scheme and a low-storage, third or fourth-order Runge-Kutta timestepping routine. It is capable of solving the compressible Navier-Stokes equations coupled with various turbulence parameterisations (e.g. Large Eddy Simulation models) and diagnostic routines. In contrast, OpenSBLI is written in Python, and by replacing the legacy core with modern code generation techniques, the existing functionality of SBLI is enriched with new flexibility; the compressible Navier-Stokes equations can still be solved in OpenSBLI for the sake of continuity, but the set of equations that can be readily solved essentially becomes a superset of that of the legacy code. Furthermore, the use of the OPS library allows the generated code to easily be targetted towards sequential, MPI, or an MPI+OpenMP hybrid backend (for CPU parallel execution), CUDA and OpenCL (for GPU parallel execution), and OpenACC (for parallel execution on accelerators), without the need to re-write the model code. OPS is readily extensible in terms of new backends, making the code generation technique an attractive way of future-proofing the codebase and preparing the framework for exascale-capable hardware when it arrives. The main achievement of OpenSBLI is the ability to express model equations at a high-level with the help of the SymPy library [17], expanding the equations based on the index notation, and coupling this functionality with the generation of OPSC-based model code and also with the OPS library which performs code targetting. OpenSBLI’s focus on the generation of computational kernels essentially forms a bridge between the high-level equations and the computational parallel loops (‘parloops’) that iterate over the grid points to solve the governing equations.

For any given simulation that is to be performed with OpenSBLI, the problem (comprising the equations to be solved, the grid to solve them on, their associated boundary and initial conditions, etc) must be defined in a setup file, which is nothing but a Python file which instantiates the various relevant components of the OpenSBLI framework. All components follow the principle of object-oriented design, and each class is explained in detail throughout the subsections that follow. An overview of the class relationships is also provided in Figure 1.

116 2.1. Equation specification

117 In a similar fashion to other problem solving environments such as Open-
 118 FOAM [18], Firedrake [4], FEniCS [5, 6], OPESCI-FD [19], Devito [20, 21],
 119 deal.II [22] and FreeFEM++ [23], OpenSBLI comprises a high-level inter-
 120 face for specifying the differential equations that are to be solved. These
 121 equations (and any accompanying formulas for temperature-dependent vis-
 122 cosity, for example) can be expressed in Einstein notation, also known as
 123 index notation. The adoption of such an abstraction is advantageous since
 124 it removes the need for the user to expand the equations by hand which can
 125 be an error-prone task. Furthermore, much like the Devito domain specific
 126 language (DSL) [20, 21] for finite difference stencil compilation, OpenSBLI
 127 makes use of the SymPy symbolic algebra library that supplies the basic com-
 128 ponents required for the modelling functionality that has been implemented
 129 in the present work. This functionality includes the automatic expansion of
 130 indices based on their contraction structure, such that repeated indices are
 131 expanded into a sum about that index, and the implementation of various
 132 types of differential operator.

133 2.1.1. Expressing

134 Consider the conservation of mass equation

$$\frac{\partial \rho}{\partial t} + \frac{\partial}{\partial x_j} [\rho u_j] = 0, \quad (1)$$

135 where u_j is the j -th component of the velocity vector \mathbf{u} , ρ is the density field,
 136 and x_j is the coordinate field in the j -th dimension. In an OpenSBLI problem
 137 setup file, the user would specify this as a string, giving the left-hand side
 138 and right-hand side of the equation in the following format:

139 `mass = "Eq(Der(rho, t), -Conservative(rho*u_j, x_j))"`

140 The functions `Der` and `Conservative` here are OpenSBLI-specific deriva-
 141 tive operators, each defined in their own class derived from SymPy's `Function`
 142 class. Other high-level interfaces such as OpenFOAM offer similar differen-
 143 tial operators such as `div` and `grad`, for example [18]. General derivatives
 144 are represented using the `Der` operator, whereas the `Conservative` oper-
 145 ator ensures that the derivative will not be expanded using the product
 146 rule. A skew-symmetric form of the derivative is also available using the
 147 `Skew` function, discussed later in Section 3.3. All of these are essentially
 148 ‘handler’/placeholder objects that OpenSBLI uses for spatial/temporal dis-
 149 cretisation after parsing and expanding the equations about the Einstein

indices. Special functions such as the Kronecker delta function and the Levi-Civita symbol are also available, derived from SymPy's `LeviCivita` and `KroneckerDelta` classes in order to handle Einstein expansion; these too are expanded later by OpenSBLI.

2.1.2. Parsing

Once all of the governing equations have been expressed by the user in string format, they are collected together in OpenSBLI's `Problem` class (see Appendix A). This class also accepts *substitutions*, *formulas*, and *constants*. For long equations, such optional *substitutions* (such as the definition of the stress tensor) can be written as a separate string (in the same way as the governing equations) to allow better equation readability, and then automatically substituted into the equations (such as the conservation of momentum and energy equations) at expansion-time instead of performing such error-prone manipulations by hand. The constitutive equations which define a relationship between the prognostic and non-prognostic variables are given as *formulas*, for example temperature-dependent viscosity relations, and an equation of state for pressure. The *constants* are the spatially and temporally independent variables which are represented as strings. Upon instantiation of the `Problem` class, the process is invoked to transform the equations into their final expanded form.

For each equation in string form, a new OpenSBLI `Equation` object is created. During its initialisation, SymPy's `parse_expr` function converts the equation string into a SymPy `Eq` data type. Any of the OpenSBLI derivative operators such as `Der` and `Conservative` (currently in string format) are replaced by actual instances of the `Der` and `Conservative` classes. Similarly, any substitutions given in the `Problem` are parsed and substituted directly into the expression using SymPy's `xreplace` function. All other terms in the parsed expression are represented by OpenSBLI's `EinsteinTerm` class, derived from SymPy's `Symbol` class, which contains its own methods and attributes for determining/expanding Einstein indices. For example, the class's initialisation method `__init__` splits up the term `u_j` where there are underscore markers, and stores the Einstein index `j` in a list as a SymPy `Idx` object. The `get_expanded` method later replaces the alphabetical Einstein indices with actual numerical indices, replacing `_j` with 0 and 1, in the 2D case. Finally, any constants in the `Problem` object are also represented as an `EinsteinTerm` object, but are flagged as constant terms in OpenSBLI, so that they are not spatially or temporally-dependent. The coordinate vector

187 components x_j (and the time term t) are a special case of an `EinsteinTerm`;
 188 these are marked with a `is_coordinate` flag so that, during the expansion
 189 phase, the `EinsteinTerms` are made dependent on the coordinate field (and
 190 time, if appropriate) to ensure that differentiation is performed correctly.

191 2.1.3. Expanding

192 After the parsing and substitution stage, the equations are expanded
 193 about repeated indices. Note that this process is performed by `OpenSBLI`,
 194 although various SymPy classes underpin the functionality. Following the
 195 example, (1) would be expanded as

$$\frac{\partial \rho}{\partial t} + \frac{\partial}{\partial x_0} [\rho u_0] + \frac{\partial}{\partial x_1} [\rho u_1] = 0. \quad (2)$$

196 `OpenSBLI` loops over each `EinsteinTerm` stored in the parsed `Equation`
 197 object, and maps it to a SymPy `Indexed` object. For example, the term `u_k`
 198 would first be mapped to `u[k]`. The index `k` in the term is then expanded
 199 over $0, \dots, d-1$ (where d is the dimension of the problem) by replacing it with
 200 each integer dimension, yielding a SymPy `MutableDenseNDimArray` array of
 201 size d (for a vector function, or $d \times d$ for a tensor of rank 2) of expanded
 202 variables which is stored as a class attribute. For example, expanding the
 203 vector `u[k]` yields the expansion array `[u0, u1]` in 2D. Upon expansion,
 204 the terms are also made spatially-dependent (i.e. indexed by x_0, x_1, x_2
 205 coordinates, depending on the dimension) and, if applicable, temporally-
 206 dependent (i.e. indexed also by t). The only exceptions to this are constants
 207 such as the Reynolds number `Re`. The expansion array from the previous
 208 example then becomes `[u0[x0, x1, t], u1[x0, x1, t]]` (and `[x0, x1]`
 209 for the constant coordinate field).

210 Each equation is expanded by locating any repeated indices and then sum-
 211 ming over them as appropriate. For example, after mapping each `EinsteinTerm`
 212 (e.g. `u_k`) to an `Indexed` object (e.g. `u[k]`), the mass equation is represented
 213 internally as

214 `Eq(Der(rho, t), -Conservative(rho*u[k], x[k]))`

215 Since the index `k` is repeated, the expansion arrays are used to expand
 216 this expression to

217 `Eq(Der(rho[x0, x1, t], t), -Conservative(rho[x0, x1, t]*u0[x0,`
 218 `x1, t], x0[x0, x1, t]) - Conservative(rho[x0, x1, t]*u1[x0, x1, t]),`
 219 `x1[x0, x1, t]))`

220 Finally, the `Der` and `Conservative` functions are applied, with the ex-
 221 pression becoming
 222 `Eq(Derivative(rho[x0, x1, t], t), -Derivative(rho[x0, x1, t]*u0[x0,`
 223 `x1, t], x0) - Derivative(rho[x0, x1, t]*u1[x0, x1, t], x1))`
 224 which is equivalent to (2). Similar expansion can also be applied for any
 225 other equations involving e.g. diagnostic fields. Note how the calls to `Der`
 226 and `Conservative` have been replaced by calls to SymPy’s `Derivative` class
 227 (which in turn uses SymPy’s `diff` function); while it is SymPy that handles
 228 the differentiation, it is OpenSBLI that handles the exact formulation of
 229 the derivative (i.e. OpenSBLI has ensured that the derivative has not been
 230 expanded using the product rule here).

231 Any nested derivatives are also handled here. It is not currently possible
 232 to specify, for example, `diff(diff(u_j, x_i), x_j)` using SymPy’s `diff`
 233 function directly because the fact that `u_j` is dependent on `x_i` and `x_j` is
 234 not taken into account. In contrast, the use of `Der` and `EinsteinTerms` like
 235 `u_j` in OpenSBLI allows the derivative to be computed correctly since the
 236 terms are made dependent through the use of `Indexed` objects as previously
 237 described. OpenSBLI users must instead use the `Der` function `Der(Der(u_j,`
 238 `x_i), x_j)`. For each nested derivative (or nested function in general), the
 239 inner function is evaluated first along with all other non-nested functions.
 240 Only then is the outer function applied.

241 For the purposes of debugging, OpenSBLI includes a `LatexWriter` class
 242 that takes the expanded equations as input and writes them out in LaTeX
 243 format so developers can more easily spot errors, for example where indices
 244 have been expanded incorrectly.

245 2.2. Grid

246 The governing equations are discretised on a regular grid of solution points
 247 that span the domain of interest; an example is provided in Figure 2. All grid-
 248 related functionality is handled by the `Grid` class, which must be instantiated
 249 by the user in the problem setup file. The dimensionality of the problem d ,
 250 the number of points in each dimension, and the grid spacing must all be
 251 supplied. A problem of dimension d would generate a grid of $N_{x_0} \times \dots \times N_{x_{d-1}}$
 252 solution points in total, where N_{x_i} represents the user-defined number of grid
 253 points in direction x_i .

254 For the sake of looping over each solution point and computing the nec-
 255 essary derivatives via the finite difference method, each (non-constant) term
 256 is processed further by OpenSBLI; the index of each spatial coordinate (e.g.

257 $\mathbf{x}0$) is mapped onto an index over the grid points in that spatial direction
 258 (e.g. $i0$) which will iterate from 0 to $N_{x_i} - 1$ (for a given direction x_i) when
 259 the computational kernel is eventually generated.

260 In addition to the solution points within the physical domain, a set of
 261 halo points (or ‘ghost’ points), which border the outer-most grid points, are
 262 also created automatically depending on the boundary conditions and the
 263 spatial order of accuracy. These halo points are necessary to ensure that
 264 the derivatives near the boundary can be computed with the same stencil
 265 as the ‘inner’ points. The exact number of halo points required therefore
 266 depends on the number of stencil points; for example, in Figure 2 the stencil
 267 for a second-order central difference (using 3 points in each direction) would
 268 require one halo point at each end of the domain. The values that these halo
 269 points hold depend on the type of boundary condition applied, and this is
 270 discussed in more detail in Section 2.6.

271 Every field/term in the governing equations that is represented by the grid
 272 indices holds a so-called ‘work array’ which essentially contains the field’s nu-
 273 merical value at each of the grid points, including the halos. The implemen-
 274 tation of initial and boundary conditions is done by accessing and modifying
 275 this work array, as will be described in Sections 2.5 and 2.6.

276 *2.3. Computational kernels*

277 The `Kernel` class defines a sequence of computational steps that should
 278 be performed to solve the governing equations. For instance, one kernel may
 279 be created to compute the spatial derivative of a field, while another kernel
 280 handles the initialisation of the field values based on a given initial condition,
 281 and another handles the enforcement of boundary conditions that involve
 282 computations. During the instantiation of a kernel, the relevant variables and
 283 fields are classified as inputs, outputs and input/outputs (i.e. both an input
 284 and an output), and the kernel’s range of evaluation (i.e. the range of grid
 285 indices over which the kernel is applied). This helps to minimise data transfer,
 286 since only those variables/fields required to perform the computation are
 287 passed to the generated kernel code.

288 *2.4. Discretisation schemes*

289 Once a grid is created, the equations are discretised upon that grid. For
 290 spatial discretisation purposes OpenSBLI offers a central differencing scheme
 291 for first and second-order derivatives; all the stencil coefficients are computed
 292 using SymPy, which allows stencils of an arbitrary order of accuracy to be

293 created. For temporal discretisation purposes, OpenSBLI features the (first-
 294 order) forward Euler scheme as well as the same low-storage, third-order
 295 Runge-Kutta timestepping scheme [24] present in the legacy SBLI code.

296 To use a particular scheme, one should instantiate a discretisation scheme
 297 derived from the generic base class called **Scheme**, which essentially stores
 298 the finite difference stencil coefficients or the weights used in a particular
 299 time-stepping scheme. Spatial and temporal schemes should be instantiated
 300 separately.

301 For the purpose of spatial discretisation, handled by the OpenSBLI **SpatialDiscretisation**
 302 class, an **Evaluations** object is created for each of the formulas, and the
 303 derivatives in the equations. Each **Evaluations** object automatically finds
 304 and stores the dependencies of a given term (e.g. $\partial(A + B)/\partial x_0$ requires
 305 the dependencies A and B). Once all the **Evaluations** have been created,
 306 they are sorted with respect to their dependencies being evaluated (e.g. if
 307 B depends on A , then A should be evaluated first). The next step involves
 308 defining the range of grid point indices over which each evaluation should
 309 be performed, and also assigning a temporary work array for each evalua-
 310 tion. All of the evaluations are then described by a **Kernel** object (see
 311 Section 2.3). It is here, while creating the kernels, that the (continuous)
 312 spatial derivatives are automatically replaced by their discrete counterparts.
 313 It should be noted that, for the evaluation of formulas, these kernels are
 314 fused together if they have no inter-dependencies to avoid race conditions
 315 when running on threaded architectures. Finally, to evaluate the residual
 316 for the purposes of temporal discretisation, the derivatives in the expanded
 317 equations (represented by an **Evaluations** object) are substituted by their
 318 temporary work arrays, and a **Kernel** is created for evaluating the residual
 319 of each equation.

320 The temporal discretisation, handled by the **TemporalDiscretisation**
 321 class, involves applying the various stages of the time-stepping scheme sup-
 322 plied using the residuals computed by the spatial discretisation process. Sim-
 323 ilarly, a **Kernel** object is created for the evaluations in the time-stepping
 324 scheme.

325 2.5. Initial conditions

326 In order for the prognostic fields to be advanced forward in time, initial
 327 conditions can be applied using the **GridBasedInitialisation** class. This
 328 is accomplished in much the same way as specifying equations, but involves

329 assignment of grid variables and work arrays of grid point values. For exam-
 330 ple, in the simulation setup file the x_0 coordinate can be defined using the
 331 grid point index and Δx_0 :

```
332 x0 = "Eq(grid.grid_variable(x0), grid.Idx[0]*grid.deltas[0])",
333 which in turn defines the initial value for each prognostic variable, by assign-
334 ing this to the array of values at each grid point (also known as the variable's
335 work array), e.g.:
336 rho = "Eq(grid.work_array(rho), 2.0*sin(x0))".
```

337 2.6. Boundary conditions

338 OpenSBLI currently comprises two types of boundary condition, imple-
 339 mented in the classes `PeriodicBoundaryCondition` and `SymmetryBoundaryCondition`.
 340 Users may apply different boundary conditions in different directions if they
 341 so wish. Periodic boundaries are defined such that, for each prognostic field
 342 ϕ , $\phi(x_0) = \phi(x_N)$ where N is the number of points in the domain. This condi-
 343 tion is achieved via the exchange of halo point data at each end of the domain.
 344 Symmetry boundary conditions enforce the condition that $\phi(x_N) = \phi(x_{N-1})$
 345 for scalar fields and $\phi_i(x_N) = -\phi_i(x_{N-1})$ for vector fields (in the direction i),
 346 which is achieved using a computational kernel.

347 2.7. Input and output

348 The state of the prognostic fields can be written to disk every n iterations
 349 as defined by the user, or only at the end of the simulation. This function-
 350 ality is handled by the `FileIO` class. OpenSBLI adopts the HDF5 format
 351 [25, 26] as it features parallel read/write capabilities and therefore has the
 352 potential to overcome the serial input/output bottleneck currently plaguing
 353 many large-scale parallel applications [27, 28]. Future releases of OpenSBLI
 354 will come with the ability to read in mesh files and the state fields from an
 355 HDF5 file, enabling the restarting of simulations from ‘checkpoints’ as well
 356 as the assignment of initial conditions that cannot be simply defined by a
 357 formula.

358 2.8. Code generation

359 OpenSBLI currently generates code in the OPSC language which per-
 360 forms the simulation; this is essentially standard C++ code that includes calls
 361 to the OPS library. Such functionality is accomplished using the OpenSBLI
 362 `OPSCCodePrinter` class (derived from SymPy’s `CCodePrinter` class, used to
 363 perform the generation of OPSC code statements) and the `OPSC` class (which

364 agglomerates the literal strings of OPSC statements and kernel functions and
 365 writes them to file). The generated code's structure follows a generic tem-
 366 plate that maps out the order in which the simulation steps/computations
 367 are to be called. The template is represented as a multi-line Python string
 368 template, with each line containing a place-holder for the code that per-
 369 forms a particular step. Examples include `$header` which is replaced by any
 370 generic boilerplate header code (e.g. `#include <stdlib.h>` and kernel func-
 371 tion prototypes), `$initialisation` which is replaced by the grid and field
 372 setup (e.g. by declaring an OPS block using the `ops_decl_block` function),
 373 and `$bc_calls` which is replaced by calls to the boundary condition kernel(s).
 374 This template can be readily changed to incorporate additional functionality,
 375 such as the inclusion of turbulence models. Once all component place-holders
 376 have been replaced by OPSC code, the code is written out to disk. For the
 377 case of the OPSC language, two files are written; one is a C++ header file
 378 containing the computational kernels, and the other is the C++ source file
 379 containing various constant definitions (e.g. the timestep size `delta_t`, and
 380 the constants of the Butcher tableau for the time-stepping scheme), OPS
 381 data structures, and calls to the kernels specified in the header file.

382 OpenSBLI's local Python objects (most pertinently, the kernel objects
 383 that describe the computations to be performed on the grid) are essentially
 384 translated to OPSC data structures and function calls during the prepara-
 385 tion of the code. For instance, when declaring computational stencils that
 386 define a particular central differencing scheme, the local grid indices stored
 387 in the `Central` scheme object are used to write out an `ops_stencil` defini-
 388 tion during code generation. Similarly, `ops_halo` structures and calls to
 389 `ops_halo_transfer` are produced to facilitate the implementation of the pe-
 390 riodic boundary conditions. All fields are declared as `ops_dat` datasets; for
 391 an example of where these are used, see the function `ops_argument_call`
 392 in the file `opsc.py` which generates/accumulates calls to the OPS function
 393 `ops_arg_dat` through the use of 'printf'-style string formatting, filling in
 394 the 'placeholder' arguments (e.g. `%s` in Python) with values from the local
 395 OpenSBLI objects. Finally, calls to OpenSBLI `Kernel` objects are repre-
 396 sented in OPSC as regular C++ functions (see Figure 3) which are passed
 397 to the `ops_par_loop` function (see Figure 4), which executes the function ef-
 398 ficiently over the range of grid points within the desired block; OpenSBLI
 399 is currently a single-block code so only one block, containing all the grid points,
 400 is used. Further details on the OPS data structures and functionality can be
 401 found in the work by [10].

Some optimisations are performed during the code generation stage by OpenSBLI to avoid unnecessary and expensive division operations in the kernels; rational numbers (e.g. finite difference stencil weights that are rational) and constant `EinsteinTerms` raised to negative powers (e.g. Re^{-1}) are evaluated and stored (e.g. by over-riding the `_print_Rational` method in the `OPSCCodePrinter` class).

Once the code generation process is complete, the OPS library is called to target the code towards various backends. These include the sequential code, MPI and hybrid MPI+OpenMP parallellised versions of the code for CPUs, CUDA and OpenCL versions of the code for GPUs, and an OpenACC version for accelerators. The test cases presented in this paper (see Section 3) consider the sequential, MPI, and CUDA backends. Targetting ‘hand-written’/manually-generated model code towards a particular architecture is something that is well-known as a time-consuming, error-prone and often unsustainable activity; often numerical models have to be completely re-written, involving many if-else statements and `#ifdef`-style pragmas to ensure that the correct branch of the code is followed for a given backend. As the number of backends grows, the code becomes unsustainable. In contrast, with the abstraction introduced here through code generation, support for a new backend only needs to be added to the OPS library; the top-level, abstract definition of the equations and their implementation need not be modified due to the separation of concerns, thereby highlighting one of the key advantages of automated model development.

When comparing the number of lines and the complexity of the code that gets generated by OpenSBLI, another advantage of automated model development becomes clear; in the case of the 3D Taylor-Green vortex test case, the problem specification file containing ~ 100 lines generates OPSC code that is approximately 1,500 lines long (excluding blank lines and comments). As more parameterisations (e.g. Large Eddy Simulation turbulence models) and diagnostic field computations are added, it is expected that this number would grow even further relative to the number of lines required in the setup file.

3. Verification and Validation

In order to verify the correctness of OpenSBLI and be confident in the ability of the solution algorithms to accurately represent the underlying

437 physics, three representative test cases covering 1, 2 and 3 dimensions were
 438 created and are presented here.

439 3.1. Propagation of a wave

440 This 1D test case considers the first-order wave equation, given by

$$\frac{\partial \phi}{\partial t} + c \frac{\partial \phi}{\partial x} = 0, \quad (3)$$

441 where ϕ is the quantity that is transported at constant speed c . The expected
 442 behaviour is that an arbitrary initial profile at time $t = 0$ is displaced by a
 443 distance $d_t = ct$, such that $\phi(x, t = 0) = \phi(x = d_t, t = T)$ for some finish
 444 time T . The constant c was set to 0.5 ms^{-1} in this case, and the equation was
 445 solved on the line $0 \leq x \leq 1 \text{ m}$. Eighth-order central differencing was used to
 446 discretise the domain in space in conjunction with a third-order Runge-Kutta
 447 scheme for temporal discretisation. The grid spacing Δx was set to 0.001 m ,
 448 and the timestep size Δt was set to $4 \times 10^{-4} \text{ s}$, yielding a Courant number
 449 of 0.2 . A smooth, periodic initial condition $\phi(x, t = 0) = \sin(2\pi x)$ was used,
 450 and periodic boundary conditions were enforced at both ends of the domain.

451 The simulation was run in serial (on an Intel® Core™ i7-4790 CPU)
 452 until a finish time of $t = 1 \text{ s}$. The initial and final states of the solution
 453 field ϕ are shown in Figure 5. As desired, the error in the solution is very
 454 small at $O(10^{-10})$, and provides some confidence in the implementation of
 455 the solution method and the periodic boundary conditions.

456 3.2. Method of manufactured solutions

457 The method of manufactured solutions (MMS) is a rigorous way to check
 458 the correctness of a numerical method's implementation [29, 30, 31]. The
 459 overall algorithm involves constructing a manufactured solution ϕ_m for the
 460 prognostic variable(s) ϕ and substituting this into the governing equation.
 461 Since the manufactured solution will not, in general, be the exact solution
 462 to the equation, a non-zero residual term will be present. This residual
 463 term is then subtracted from the RHS such that the manufactured solution
 464 essentially becomes the exact/analytical solution of the modified equation
 465 (i.e. the one with the source term). A suite of simulations can then be
 466 performed using increasingly fine grids to check that the numerical solution
 467 converges to the manufactured solution at the expected rate determined by
 468 the discretisation scheme.

469 For this test, the 2D advection-diffusion equation (with a source term S)
 470 given by

$$\frac{\partial \phi}{\partial t} + \frac{\partial}{\partial x_j} \left[\phi u_j - k \frac{\partial \phi}{\partial x_j} \right] + S = 0, \quad (4)$$

471 is considered.

472 The constant k is the diffusivity coefficient which is set to $0.75 \text{ m}^2\text{s}^{-1}$
 473 here. The prescribed field u_i is the i -th velocity component, with $u_0 = 1.0$
 474 ms^{-1} and $u_1 = -0.5 \text{ ms}^{-1}$. The prognostic field ϕ is to be determined and
 475 has an initial condition of $\phi(x, t = 0) = 0$. In a similar fashion to the works
 476 of [29, 30, 31], the manufactured/‘analytical’ solution $\phi_m = \sin(x_0) \cos(x_1)$
 477 employs a mixture of sine and cosine functions since these are continuous and
 478 infinitely differentiable. The SAGE framework [32] was used to symbolically
 479 determine the residual/source term S .

480 The domain is a 2D square with dimensions $0 \leq x_0 \leq 2\pi \text{ m}$ and $0 \leq$
 481 $x_1 \leq 2\pi \text{ m}$ such that the manufactured solution is periodic. Furthermore,
 482 periodic boundary conditions are applied on all sides of the domain. Six
 483 central differencing schemes of order 2, 4, 6, 8, 10 and 12 are considered
 484 for the spatial discretisation, and a third-order Runge-Kutta scheme is used
 485 throughout to advance the equation in time. To perform the convergence
 486 analysis, the grid spacing was halved for each successive case such that Δx
 487 $= \Delta y = \frac{\pi}{2}, \frac{\pi}{4}, \frac{\pi}{8}, \frac{\pi}{16}$ and $\frac{\pi}{32}$. The timestep size Δt was also halved for each
 488 case to maintain a maximum bound of 0.025 on the Courant number; this
 489 was purposefully kept small and near-constant to minimise the influence of
 490 temporal discretisation error [33]. All simulations were run in serial (on an
 491 Intel® Core™ i7-4790 CPU) until a finish time of $T = 100 \text{ s}$ to ensure that
 492 a steady-state solution was attained.

493 Figure 6 demonstrates how ϕ converges towards the manufactured so-
 494 lution ϕ_m as the grid is refined. The convergence rate for each order of
 495 the central difference scheme is illustrated in Figure 7. The anomaly in the
 496 twelfth-order convergence plot was likely caused by reaching the limit of ma-
 497 chine precision. Overall, these results provide confidence in the correctness
 498 of the automatically-generated code/model.

499 3.3. 3D Taylor-Green vortex

500 The Taylor-Green vortex is a well-known hydrodynamic problem [34, 35,
 501 36] characterised by transition to turbulence, decay of turbulence, and the
 502 energy dissipation during its evolution. It is frequently used to evaluate the

503 ability of a numerical method to capture the underlying physical processes.
 504 During the initial stages of evolution, the dynamics display structural changes
 505 (rolling up, stretching and interaction of the vortices). This process is inviscid
 506 in nature. Later the vortices break down and transition into fully-turbulent
 507 dynamics. As there are no external forces or turbulence-generating mecha-
 508 nisms, the small-scale structures dissipate all the energy, and the fluid even-
 509 tually comes to rest [34]. The numerical method employed should be able to
 510 capture each of these stages accurately.

511 The 3D compressible Navier-Stokes equations were solved in non-dimensional
 512 form, written in Einstein notation as

$$\frac{\partial \rho}{\partial t} + \frac{\partial}{\partial x_j} [\rho u_j] = 0, \quad (5)$$

$$\frac{\partial \rho u_i}{\partial t} + \frac{\partial}{\partial x_j} [\rho u_i u_j + p \delta_{ij} - \tau_{ij}] = 0, \quad (6)$$

513 and

$$\frac{\partial \rho E}{\partial t} + \frac{\partial}{\partial x_j} [\rho E u_j + u_j p - q_j - u_i \tau_{ij}] = 0. \quad (7)$$

514 for the conservation of mass, momentum and energy, respectively. The (di-
 515 mensionless) quantity ρ is the fluid density, u_i is the i -th (scalar) component
 516 of the velocity vector \mathbf{u} , p is the pressure field, E is the total energy. The
 517 components of the stress tensor τ are given by

$$\tau_{ij} = \frac{1}{\text{Re}} \left(\frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i} - \frac{2}{3} \delta_{ij} \frac{\partial u_k}{\partial x_k} \right), \quad (8)$$

518 where δ_{ij} is the Kronecker Delta function and Re is the Reynolds number.
 519 The components of the heat flux term q are given by

$$q_j = \frac{\mu}{(\gamma - 1) \text{M}^2 \text{Pr} \text{Re}} \frac{\partial T}{\partial x_j}, \quad (9)$$

520 where T is the temperature field, γ is the ratio of specific heats, M is the
 521 Mach number, and Pr is the Prandtl number. The various quantities are
 522 non-dimensionalised using the reference velocity u_{ref} , the reference length L ,
 523 the reference density ρ_{ref} , and the reference temperature T_{ref} .

524 The equation of state linking p , ρ and T , is defined by

$$p = \frac{1}{\gamma M^2} \rho T, \quad (10)$$

525 and the total energy is given by

$$\rho E = \frac{p}{\gamma - 1} + \frac{1}{2} \rho u_j^2. \quad (11)$$

526 The pressure p is non-dimensionalised by $\rho_{\text{ref}} u_{\text{ref}}^2$.

527 Central finite difference schemes are non-dissipative and are therefore
 528 suitable for accurately capturing turbulent dynamics. However, the lack of
 529 dissipation can make the scheme unstable. To improve the stability, a skew-
 530 symmetric formulation [37, 38, 39, 40] was applied to the convective terms
 531 in (5), (6) and (7); the convective term then becomes

$$\frac{\partial}{\partial x_j} [\rho \phi u_j] = \frac{1}{2} \left(\frac{\partial}{\partial x_j} \rho \phi u_j + u_j \frac{\partial}{\partial x_j} \rho \phi + \rho \phi \frac{\partial}{\partial x_j} u_j \right), \quad (12)$$

532 where ϕ should be set to 1, u_j and E for the continuity, momentum and
 533 energy equations, respectively. It should also be noted that the both the
 534 convective and viscous terms are discretised using the same spatial order.
 535 In all of the simulations performed, the Laplacian in the viscous term is
 536 expanded using a finite difference representation of the second derivative
 537 (i.e. not treated by successive first derivatives).

538 As per the work of [35] and [36], the equations were solved in a 3D cube,
 539 with $0 \leq x_0 \leq 2\pi L$, $0 \leq x_1 \leq 2\pi L$, and $0 \leq x_2 \leq 2\pi L$. Periodic boundary
 540 conditions were applied on all surfaces. The following initial conditions were
 541 imposed at time $t = 0$:

$$u_0(x_0, x_1, x_2, t = 0) = \sin\left(\frac{x_0}{L}\right) \cos\left(\frac{x_1}{L}\right) \cos\left(\frac{x_2}{L}\right), \quad (13)$$

$$u_1(x_0, x_1, x_2, t = 0) = -\cos\left(\frac{x_0}{L}\right) \sin\left(\frac{x_1}{L}\right) \cos\left(\frac{x_2}{L}\right), \quad (14)$$

$$u_2(x_0, x_1, x_2, t = 0) = 0, \quad (15)$$

$$p(x_0, x_1, x_2, t = 0) = \frac{1}{\gamma M^2} + \frac{1}{16} \left(\cos\left(\frac{2x_0}{L}\right) + \cos\left(\frac{2x_1}{L}\right) \right) \left(2 + \cos\left(\frac{2x_2}{L}\right) \right), \quad (16)$$

542 In all the simulations, $\text{Re} = 1,600$, $\text{Pr} = 0.71$, $M = 0.1$, and $\gamma = 1.4$. The
 543 reference quantities L , u_{ref} and ρ_{ref} were set to 1.0, and the reference tem-
 544 perature T_{ref} was evaluated using the equation of state (10).

545 A fourth-order accurate central differencing scheme was used to spatially
 546 discretise the domain, and a third-order Runge-Kutta timestepping scheme
 547 was used to march the equations forward in time. A set of simulations
 548 was performed over a range of resolutions, namely 64^3 , 128^3 , 256^3 and 512^3
 549 uniformly-spaced grid points. For the 64^3 case, a non-dimensional time-
 550 step size Δt of 3.385×10^{-3} [35] was used. Each time the number of grid
 551 points was doubled, the time-step size was halved to maintain a constant
 552 upper bound on the Courant number. The generated code was targetted
 553 towards the CUDA backend using OPS and executed on an NVIDIA Tesla
 554 K40 GPU until a non-dimensional time of $t = 20$, except for the 512^3 case;
 555 this was targetted towards the MPI backend and run in parallel over 1,440
 556 processes on the UK National Supercomputing Service (ARCHER) due to
 557 lack of available memory on the GPU, and provided a good example of how
 558 the backend can be readily changed.

559 The z -component of the vorticity field at various times can be found in
 560 Figure 8. At non-dimensional time $t = 2.5$ vortex evolution and stretching
 561 are clearly visible, progressing onto highly turbulent dynamics where the
 562 relatively smooth structures roll-up and eventually breakdown at around t
 563 $= 9$. This point is characterised by peak enstrophy in the system. The final
 564 stage of the simulation features the decay of the turbulent structures such
 565 that the enstrophy tends towards its initial value.

566 Following the definitions of [35], the integrals of the kinetic energy

$$E_k = \frac{1}{\rho_{\text{ref}}\Omega} \int_{\Omega} \frac{1}{2} \rho u_j u_j \, d\Omega, \quad (17)$$

567 and enstrophy

$$\varepsilon = \frac{1}{\rho_{\text{ref}}\Omega} \int_{\Omega} \frac{1}{2} \rho \left(\epsilon_{ijk} \frac{\partial u_k}{\partial x_j} \right)^2 \, d\Omega, \quad (18)$$

568 were computed throughout the simulations. Note that Ω is the whole domain
 569 and ϵ_{ijk} is the Levi-Civita function. These quantities are shown in Figures 9
 570 and 10 for the various grid resolutions, and are plotted against the reference
 571 data from a spectral element simulation by [41] using a 512^3 grid for com-
 572 parison. Figure 10 highlights the inviscid nature of the Taylor-Green vortex
 573 problem for $t < \sim 3$ -4. The transition to turbulence occurs from $\sim 3 < t < 9$
 574 (which is associated with the peak in enstrophy in Figure 9). Finally, dissipa-
 575 tion occurs at $t > 9$. The results show a clear agreement with the reference

576 data, and represents a solid first step towards the validation of OpenSBLI.

577 4. Conclusion

578 Advances in compute hardware are driving a need to change the current
579 state of numerical model development. By developing a new modelling frame-
580 work based on automated solution techniques, we have effectively future-
581 proofed the core of the SBLI codebase; no longer does a computational sci-
582 entist need to re-write significant portions of code in order to get it up and
583 running on a new piece of hardware. Instead, the model is derived from a
584 high-level specification independent of the architecture that it will run on,
585 and the underlying code is automatically generated and tailored to a particu-
586 lar backend, the responsibility for which would rest with computer scientists
587 who are experts in parallel programming paradigms. Furthermore, the ease
588 at which the governing equations can be changed is a fundamental advantage
589 of using such abstract specifications. This was highlighted here by consider-
590 ing three test cases, each of which comprised a different set of equations. The
591 discretisation, code generation and code targetting is performed automati-
592 cally, thereby reducing development costs and potentially avoiding errors,
593 bugs, and non-performant/non-optimal operations. In addition, code that
594 solves the different variants of the same governing equations can be easily
595 generated. For example, in the compressible Navier-Stokes equations, vis-
596 cosity can be treated either as a constant or as a spatially-varying term.
597 In static, hand-written codes this flexibility comes at the cost of writing
598 different routines for the various formulations, unlike with automated code
599 generation techniques. This is particularly useful when wanting to switch be-
600 tween Cartesian and generalised coordinates. This particular framework also
601 facilitates the fast and efficient switching between different spatial orders of
602 accuracy, and reduces the development time and effort when wishing to try
603 out new numerical formulations of the equations (or a new spatial/temporal
604 scheme) on a wide variety of test cases.

605 4.1. Future work

606 Explicit schemes such as the one implemented here can be readily ex-
607 tendible to a range of application areas such as computational aeroacoustics,
608 aero-thermodynamics, problems involving shocks, and hypersonic flow. In-
609 compressible flows may also be handled with the explicit, compressible solver
610 in OpenSBLI so long as the Mach number is sufficiently small. However, this

611 puts tight restrictions on the time-step size thereby limiting the efficiency
612 of the solver, and thus limits the range of applications that OpenSBLI can
613 handle within the context of CFD in general.

614 Extending to implicit timestepping schemes requires backend support
615 from OPS for matrix inversion, for example. Once this support is imple-
616 mented, extending OpenSBLI to handle implicit timestepping schemes is
617 straightforward.

618 The treatment of incompressible flows can be accomplished using schemes
619 such as pressure projection methods [42, 43]. Such a method requires (1) the
620 solution of a tentative velocity field, (2) the solution to a pressure Poisson
621 equation (using either direct or iterative solvers), and (3) the update/correction
622 of the velocity field. The equations defining each step would need to be given
623 by the user in OpenSBLI, in a similar fashion to implementing a projection
624 method in the Unified Form Language (UFL) in FEniCS [44, 7]. OpenS-
625 BLI would also need to recognise that a projection method has been chosen,
626 possibly via a flag set in the problem definition file. For step (2) of the
627 method, direct solvers can be implemented directly once support for ma-
628 trix inversion and fast Fourier transforms (for example) are included in OPS
629 (which in turn would need to link to various linear algebra packages such as
630 PETSc [45]). This is similar to how an implicit time-stepping scheme would
631 be implemented in OpenSBLI. On the other hand, explicit solution schemes
632 require an iterative solution to the pressure Poisson equation; this is possible
633 by writing the relevant kernel support with an exit criterion (which exits the
634 kernel once a desired tolerance for the solution residual has been attained)
635 and modifying how the code is generated with this in mind. Other methods
636 for incompressible flows such as artificial compressibility methods [42, 46, 47]
637 can be implemented with the current functionality by modifying the input
638 equations in the problem setup file accordingly.

639 The work considered here only focussed on the MPI and CUDA back-
640 ends for CPU and GPU execution, respectively. Future work will consider
641 the CPU and GPU performance on other backends, such as OpenMP. For
642 problems such as Mandelbrot Set computation and matrix multiplication,
643 OpenMP has been demonstrated to perform well against other APIs such
644 as CUDA and OpenACC [48]. Future work will also look at comparing the
645 performance of the legacy Fortran-based SBLI code against the OpenSBLI-
646 generated code in order to demonstrate the potential speed-ups that can be
647 obtained.

648 5. Code Availability

649 OpenSBLI is an open-source release of the original SBLI code developed
650 at the University of Southampton, and is available under the GNU General
651 Public Licence (version 3). Prospective users can download the source code
652 from the project’s Git repository (online location to be confirmed), or from
653 Figshare: <https://figshare.com/s/c768e8b8af4af3615341>

654 6. Acknowledgments

655 CTJ was supported by a European Union Horizon 2020 project grant en-
656 titled “ExaFLOW: Enabling Exascale Fluid Dynamics Simulations” (grant
657 reference 671571). SPJ was supported by an EPSRC grant entitled “Future-
658 proof massively-parallel execution of multi-block applications” (EP/K038567/1).
659 The data behind the results presented in this paper will be available from
660 the University of Southampton’s institutional repository. The authors ac-
661 knowledge the use of the UK National Supercomputing Service (ARCHER),
662 with computing time provided by the UK Turbulence Consortium (EPSRC
663 grant EP/L000261/1). The authors would also like to thank the NVIDIA
664 Corporation for donating the Tesla K40 GPU used throughout this research.

665 Appendix A. Example of a simulation setup file

666 The code in Figure A.11 contains the key components of a simulation
667 setup file. Specifically, this is taken from the Taylor-Green vortex simulation.
668 Other examples of setup files can be found in the `apps` directory of OpenSBLI.

669 References

- 670 [1] P. Thibodeau, Scientists, IT Community Await Exascale Computers
671 (2009).
672 URL <http://www.computerworld.com/article/2550451/computer-hardware/scientists>
- 673 [2] F. Rathgeber, G. R. Markall, L. Mitchell, N. Lorient, D. A. Ham,
674 C. Bertolli, P. H. Kelly, PyOP2: A High-Level Framework for
675 Performance-Portable Simulations on Unstructured Meshes, in: High
676 Performance Computing, Networking Storage and Analysis, SC Com-
677 panion, IEEE Computer Society, 2012, pp. 1116–1123.

- 678 [3] C. T. Jacobs, M. D. Piggott, Firedrake-Fluids v0.1: numerical modelling
679 of shallow water flows using an automated solution framework, *Geosci-*
680 *entific Model Development* 8 (3) (2015) 533–547. doi:10.5194/gmd-8-
681 533-2015.
- 682 [4] F. Rathgeber, D. A. Ham, L. Mitchell, M. Lange, F. Luporini, A. T. T.
683 McRae, G.-T. Bercea, G. R. Markall, P. H. J. Kelly, Firedrake: au-
684 tomatizing the finite element method by composing abstractions, *ACM*
685 *Transactions on Mathematical Software*.
686 URL <http://arxiv.org/abs/1501.01809>
- 687 [5] A. Logg, G. N. Wells, DOLFIN: Automated finite element com-
688 puting, *ACM Transactions on Mathematical Software* 37 (2).
689 doi:10.1145/1731022.1731030.
- 690 [6] A. Logg, K.-A. Mardal, G. N. Wells, et al., *Automated Solution of*
691 *Differential Equations by the Finite Element Method*, Springer, 2012.
692 doi:10.1007/978-3-642-23099-8.
- 693 [7] M. S. Alnæs, A. Logg, K. B. Ølgaard, M. E. Rognes, G. N. Wells, Uni-
694 fied Form Language: A domain-specific language for weak formulations
695 of partial differential equations, *ACM Transactions on Mathematical*
696 *Software* 40 (2).
- 697 [8] F. Luporini, A. L. Varbanescu, F. Rathgeber, G.-T. Bercea, J. Ramanu-
698 jam, D. A. Ham, P. H. J. Kelly, Cross-Loop Optimization of Arithmetic
699 Intensity for Finite Element Local Assembly, *ACM Transactions on Ar-*
700 *chitecture and Code Optimization* 11 (4). doi:10.1145/2687415.
- 701 [9] M. Giles, I. Reguly, G. Mudalige, *OPS C++ User’s Manual*, University
702 of Oxford (2015).
703 URL <http://www.oerc.ox.ac.uk/projects/ops>
- 704 [10] I. Z. Reguly, G. R. Mudalige, M. B. Giles, D. Curran, S. McIntosh-
705 Smith, The OPS Domain Specific Abstraction for Multi-Block Struc-
706 tured Grid Computations, in: *Proceedings of the 2014 Fourth Interna-*
707 *tional Workshop on Domain-Specific Languages and High-Level Frame-*
708 *works for High Performance Computing*, IEEE Computer Society, 2014,
709 pp. 58–67. doi:10.1109/WOLFHPC.2014.7.

- [11] G. R. Mudalige, I. Z. Regulý, M. B. Giles, A. C. Mallinson, W. P. Gaudin, J. A. Herdman, Performance Analysis of a High-level Abstractions-based Hydrocode on Future Computing Systems, in: Proceedings of the 5th international workshop on Performance Modeling, Benchmarking and Simulation of High Performance Computing Systems (PMBS '14). Held in conjunction with IEEE/ACM Supercomputing 2014 (SC'14), 2014.
- [12] S. P. Jammy, G. R. Mudalige, I. Z. Regulý, N. D. Sandham, M. Giles, Block structured compressible navier stokes solution using the ops high-level abstraction, in: Proceedings of the 27th International Conference on Parallel Computational Fluid Dynamics (Parallel CFD 2015), 2015.
- [13] E. Toubert, N. D. Sandham, Large-eddy simulation of low-frequency unsteadiness in a turbulent shock-induced separation bubble, *Theoretical and Computational Fluid Dynamics* 23 (2) (2009) 79–107. doi:10.1007/s00162-009-0103-z.
- [14] N. De Tullio, N. D. Sandham, Direct numerical simulation of breakdown to turbulence in a Mach 6 boundary layer over a porous surface, *Physics of Fluids* 22 (9). doi:10.1063/1.3481147.
- [15] J. Redford, N. D. Sandham, G. T. Roberts, Numerical simulations of turbulent spots in supersonic boundary layers: Effects of Mach number and wall temperature, *Progress in Aerospace Sciences* 52 (2012) 67–79. doi:10.1016/j.paerosci.2011.08.002.
- [16] B. Wang, N. D. Sandham, W. Hu, W. Liu, Numerical study of oblique shock-wave/boundary-layer interaction considering side-wall effects, *Journal of Fluid Mechanics* 767 (2015) 526–561. doi:10.1017/jfm.2015.58.
- [17] SymPy Development Team, SymPy: Python library for symbolic mathematics (2016).
URL <http://www.sympy.org>
- [18] OpenFOAM, OpenFOAM User Guide, Version 2.3.1 (2014).
- [19] T. Sun, OPESCI-FD: Automatic Code Generation Package for Finite Difference Models, Master's thesis, Imperial College London (2016).
URL <https://arxiv.org/abs/1605.06381>

- [20] N. Kukreja, M. Louboutin, F. Vieira, F. Luporini, M. Lange, G. Gorman, Devito: automated fast finite difference computation, in: In Proceedings of the Sixth International Workshop on Domain-Specific Languages and High-Level Frameworks for High Performance Computing, 2016.
URL <https://arxiv.org/abs/1608.08658>
- [21] M. Lange, N. Kukreja, M. Louboutin, F. Luporini, F. Vieira, V. Pandolfo, P. Valesko, P. Kazakas, G. Gorman, Devito: Towards a generic Finite Difference DSL using Symbolic Python, in: In Proceedings of the PyHPC 2016 Conference, 2016.
URL <https://arxiv.org/abs/1609.03361>
- [22] W. Bangerth, R. Hartmann, G. Kanschat, deal.II—A general-purpose object-oriented finite element library, *ACM Transactions on Mathematical Software* 33 (4). doi:10.1145/1268776.1268779.
- [23] F. Hecht, New development in FreeFem++, *Journal of Numerical Mathematics* 20 (3-4) (2012) 251–265. doi:10.1515/jnum-2012-0013.
- [24] M. H. Carpenter, C. A. Kennedy, Fourth-Order 2N-Storage Runge-Kutta Schemes, NASA Technical Memorandum 109112, National Aeronautics and Space Administration, Langley Research Center (1994).
- [25] M. Folk, E. Pourmal, Balancing Performance and Preservation Lessons Learned with HDF5, in: Proceedings of the 2010 Roadmap for Digital Preservation Interoperability Framework Workshop, US-DPIF '10, 2010, pp. 11:1–11:8. doi:10.1145/2039274.2039285.
- [26] A. Collette, *Python and HDF5: Unlocking Scientific Data*, O'Reilly Media, 2013. doi:10.1007/978-3-642-23099-8.
- [27] D. Padua, *Encyclopedia of Parallel Computing*, Vol. 4, Springer US, 2011. doi:10.1007/978-3-642-23099-8.
- [28] J. Fu, M. Misun, R. Latham, C. D. Carothers, Parallel I/O Performance for Application-Level Checkpointing on the Blue Gene/P System, in: Proceedings of the 2011 IEEE International Conference on Cluster Computing, 2011, pp. 465–473.

- 774 [29] K. Salari, P. Knupp, Code verification by the method of manufactured
775 solutions, Tech. Rep. SAND2000-1444, Sandia National Laboratories
776 (2000).
- 777 [30] P. J. Roache, Code Verification by the Method of Manufactured
778 Solutions, *Journal of Fluids Engineering* 124 (1) (2002) 4–10.
779 doi:10.1115/1.1436090.
- 780 [31] C. J. Roy, Review of code and solution verification procedures for com-
781 putational simulation, *Journal of Computational Physics* 205 (1) (2005)
782 131–156. doi:10.1016/j.jcp.2004.10.036.
- 783 [32] W. Stein, D. Joyner, SAGE: System for Algebra and Geometry Experi-
784 mentation, *ACM SIGSAM Bulletin* 39 (2).
- 785 [33] J. M. Vedovoto, A. da Silveira Neto, A. Mura, L. F. F. da Silva, Ap-
786 plication of the method of manufactured solutions to the verification of
787 a pressure-based finite-volume numerical scheme, *Computers & Fluids*
788 51 (1) (2011) 85 – 99. doi:10.1016/j.compfluid.2011.07.014.
- 789 [34] M. E. Brachet, D. I. Meiron, S. A. Orszag, B. G. Nickel, R. H. Morf,
790 U. Frisch, Small-scale structure of the Taylor-Green vortex, *Journal of*
791 *Fluid Mechanics* 130 (1983) 411–452. doi:10.1017/S0022112083001159.
- 792 [35] J. DeBonis, Solutions of the Taylor-Green Vortex Problem Using High-
793 Resolution Explicit Finite Difference Methods, in: 51st AIAA Aerospace
794 Sciences Meeting including the New Horizons Forum and Aerospace Ex-
795 position, Aerospace Sciences Meetings, 2013. doi:10.2514/6.2013-382.
- 796 [36] J. R. Bull, A. Jameson, Simulation of the Compressible Taylor
797 Green Vortex using High-Order Flux Reconstruction Schemes, in: 7th
798 AIAA Theoretical Fluid Mechanics Conference, AIAA Aviation, 2014.
799 doi:10.2514/6.2014-3210.
- 800 [37] G. A. Blaisdell, N. N. Mansour, W. C. Reynolds, Numerical simulations
801 of homogeneous compressible turbulence, Report TF-50, Department of
802 Mechanical Engineering, Stanford University (1991).
- 803 [38] G. A. Blaisdell, N. N. Mansour, W. C. Reynolds, Compressibil-
804 ity effects on the growth and structure of homogeneous turbu-

- 805 lent shear flow, *Journal of Fluid Mechanics* 256 (1993) 443–485.
806 doi:10.1017/S0022112093002848.
- 807 [39] G. A. Blaisdell, E. T. Spyropoulos, J. H. Qin, The effect of the for-
808 mulation of nonlinear terms on aliasing errors in spectral methods, *Ap-
809 plied Numerical Mathematics* 21 (3) (1996) 207 – 219. doi:10.1016/0168-
810 9274(96)00005-0.
- 811 [40] S. Pirozzoli, Numerical Methods for High-Speed Flows, *Annual Review*
812 of Fluid Mechanics 43 (1) (2011) 163–194. doi:10.1146/annurev-fluid-
813 122109-160718.
- 814 [41] Z. Wang, K. Fidkowski, R. Abgrall, F. Bassi, D. Caraeni, A. Cary,
815 H. Deconinck, R. Hartmann, K. Hillewaert, H. Huynh, N. Kroll, G. May,
816 P.-O. Persson, B. van Leer, M. Visbal, High-order cfd methods: current
817 status and perspective, *International Journal for Numerical Methods in*
818 *Fluids* 72 (8) (2013) 811–845. doi:10.1002/fld.3767.
- 819 [42] A. Chorin, The numerical solution of the Navier-Stokes equations for
820 an incompressible fluid, *Bulletin of the American Mathematical Society*
821 73 (6) (1967) 928–931.
- 822 [43] A. Chorin, Numerical Solution of the Navier-Stokes Equations, *Mathe-*
823 *matics of Computation* 22 (104) (1968) 745–762.
- 824 [44] The FEniCS Project, Incompressible Navier-Stokes equations, Online.
825 URL <https://fenicsproject.org/documentation/dolfin/1.0.1/python/demo/pde/navier-stokes/>
- 826 [45] S. Balay, S. Abhyankar, M. Adams, J. Brown, P. Brune, K. Buschel-
827 man, V. Eijkhout, W. Gropp, D. Kaushik, M. Knepley, L. Curfman
828 McInnes, K. Rupp, B. Smith, H. Zhang, PETSc Users Manual, Tech.
829 Rep. Revision 3.5, Argonne National Laboratory (2014).
- 830 [46] A. Chorin, A numerical method for solving incompressible viscous
831 flow problems, *Journal of Computational Physics* 2 (1) (1967) 12–26.
832 doi:10.1016/0021-9991(67)90037-X.
- 833 [47] T. Ohwada, P. Asinari, Artificial compressibility method revisited:
834 Asymptotic numerical method for incompressible navier-stokes equa-
835 tions, *Journal of Computational Physics* 229 (5) (2010) 1698–1723.
836 doi:10.1016/j.jcp.2009.11.003.

- 837 [48] C. L. Ledur, C. M. D. Zeve, J. C. S. dos Anjos, Comparative Anal-
838 ysis of OpenACC, OpenMP and CUDA using Sequential and Parallel
839 Algorithms, in: Proceedings of the 11th Workshop on Parallel and Dis-
840 tributed Processing (WSPPD), 2013, 2013.

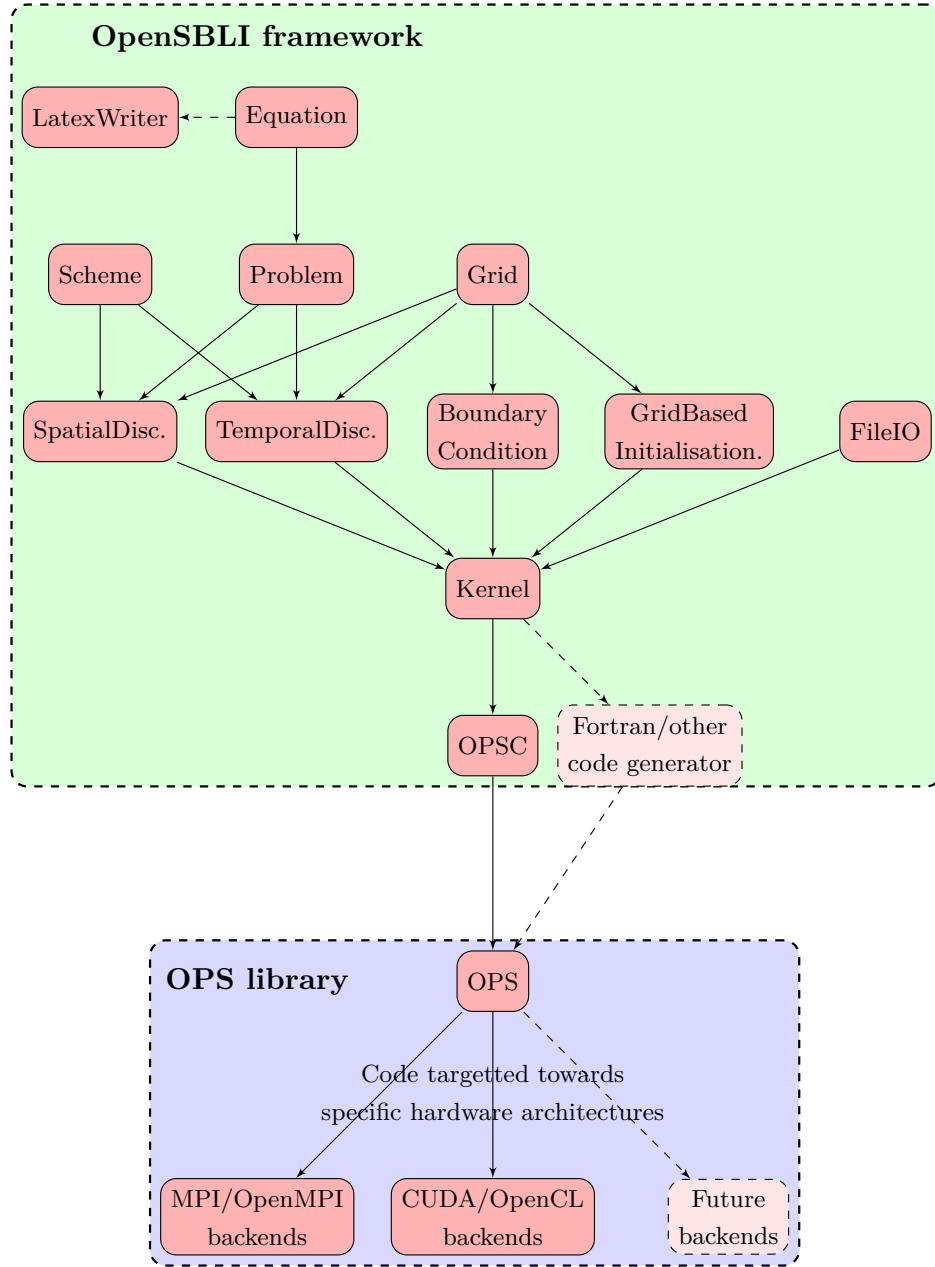


Figure 1: The overall design of the OpenSBLI framework with respect to the core classes. The code targetting happens within the OPS library. The CPU backends include MPI, OpenMP, hybrid MPI+OpenMP, as well as a sequential version of the code. The GPU backends include CUDA and OpenCL, which can also be combined with MPI to run the code on multiple GPUs in parallel. The only accelerator backend available is OpenACC.




Figure 2: The regular grid of solution points upon which the governing equations are solved. The grid point indices in the x and y directions are denoted i_0 and i_1 , respectively. The halo points that surround the outer-most points of the domain are labelled ‘H’. A computational stencil used for second-order central differencing is highlighted red in the center, with the relative grid coordinates of each point.

```

#ifdef block_0_KERNEL_H
#define block_0_KERNEL_H

void mms_2_0_block0_0_kernel(const double *phi, double *wk0)
{
    wk0[OPS_ACC1(0,0)] = (-2*phi[OPS_ACC0(0,0)] + phi[OPS_ACC0(0,-1)] + phi[OPS_ACC0(0,1)])/pow(delta1, 2);
}

void mms_2_0_block0_1_kernel(const double *phi, double *wk1)
{
    wk1[OPS_ACC1(0,0)] = (-rc0*phi[OPS_ACC0(0,-1)] + (rc0)*phi[OPS_ACC0(0,1)])/delta1;
}

```

Figure 3: Code snippet showing two kernels from a 2D ‘method of manufactured solutions’ (MMS) simulation (see Section 3.2) using second-order central differences. The first kernel computes $\frac{\partial^2 \phi}{\partial x_1^2}$ and stores it in a new work array called `wk0`. Similarly, the second kernel computes the first derivative $\frac{\partial \phi}{\partial x_1}$. The constant `delta1` represents the grid spacing in the x_1 direction, and `rc0` holds the constant value of 0.5. Calls to `OPS_ACC` are used to access the finite difference stencil structure.

```

for (int iteration=0; iteration<5093; iteration++){

    int iter_range7[] = {-1, nx0 + 1, -1, nx1 + 1};
    ops_par_loop_mms_2_0_block0_7_kernel("Save equations", mms_2_0_block, 2, iter_range7,
        ops_arg_dat(phi, 1, stencil1, "double", OPS_READ),
        ops_arg_dat(phi_old, 1, stencil1, "double", OPS_WRITE));

    for (int stage=0; stage<3; stage++){

        int iter_range0[] = {0, nx0, 0, nx1};
        ops_par_loop_mms_2_0_block0_0_kernel("D(phi[x0 x1 t] x1 x1)", mms_2_0_block, 2, iter_range0,
            ops_arg_dat(phi, 1, stencil0, "double", OPS_READ),
            ops_arg_dat(wk0, 1, stencil1, "double", OPS_WRITE));

        int iter_range1[] = {0, nx0, 0, nx1};
        ops_par_loop_mms_2_0_block0_1_kernel("D(phi[x0 x1 t] x1)", mms_2_0_block, 2, iter_range1,
            ops_arg_dat(phi, 1, stencil2, "double", OPS_READ),
            ops_arg_dat(wk1, 1, stencil1, "double", OPS_WRITE));
    }
}

```

Figure 4: Code snippet showing the calls to the two kernels in Figure 3 inside the inner for-loop. The outer loop is the time-stepping loop which iterates for a user-defined number of iterations. The inner loop iterates over the 3-stages of the low-storage, third-order Runge-Kutta scheme. Note that each kernel is actually executed as an **ops_par_loop** over all the grid points.

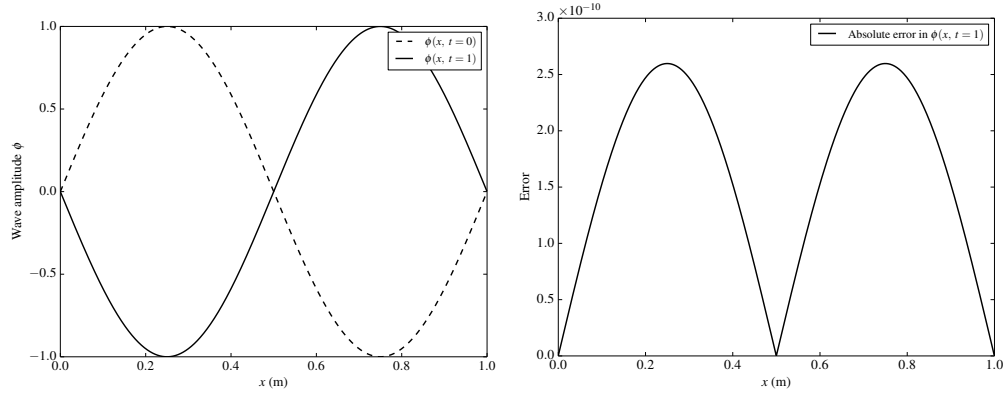


Figure 5: Results from the 1D wave propagation simulation. Left: The solution field ϕ at time $t = 0$ s and $t = 1$ s. Right: The error between the analytical solution and the numerical solution at time $t = 1$ s.




Figure 6: (a-e): The numerical solution field ϕ at the finish time $t = T$. (f): The manufactured solution ϕ_m . All results are from the twelfth-order MMS simulation set.




Figure 7: The absolute error (in the L2 norm) between the numerical solution ϕ and the exact/manufactured solution ϕ_m , from the suite of MMS simulations. The solid lines represent the expected convergence rate for each order.

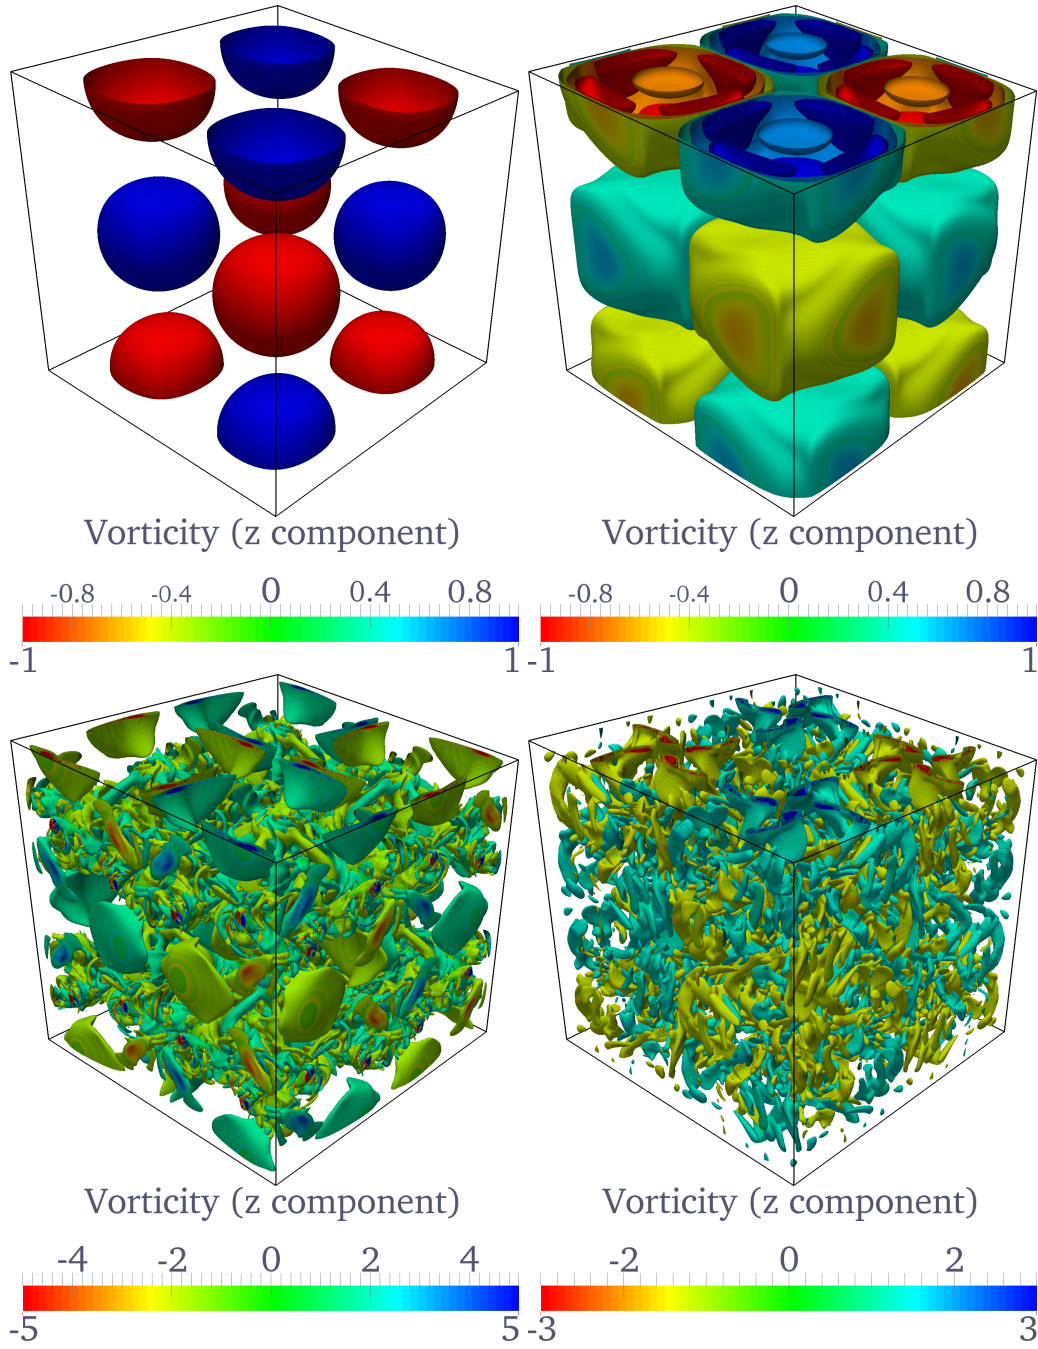


Figure 8: Visualisations of the non-dimensional vorticity (z -component) iso-contours, from the Taylor-Green vortex test case with a 256^3 grid, at various non-dimensional times. Top left to bottom right: non-dimensional time $t = 0, 2.5, 10, 20$.




Figure 9: The integral of the enstrophy in the domain until non-dimensional time $t = 20$, from the Taylor-Green vortex test case. The reference data from [41] is also shown for comparison.




Figure 10: The integral of the kinetic energy in the domain until non-dimensional time $t = 20$, from the Taylor-Green vortex test case. The reference data from [41] is also shown for comparison.

```

# Problem dimension
ndim = 3

# Define the compressible Navier-Stokes equations in Einstein notation.
mass = "Eq(Der(rho,t), - Skew(rho*u_j,x_j))"
momentum = "Eq(Der(rhou_i,t), -Skew(rhou_i*u_j,x_j) - Der(p,x_i) + Der(tau_i_j,x_j) )"
energy = "Eq(Der(rhoE,t), - Skew(rhoE*u_j,x_j) - Conservative(p*u_j,x_j) + Der(q_j,x_j) + Der(u_i*tau_i_j,x_j) )"
equations = [mass, momentum, energy]

# Substitutions and constants
stress_tensor = "Eq(tau_i_j, (1.0/Re)*(Der(u_i,x_j) + Der(u_j,x_i)- (2/3)*KD(i,j)*Der(u_k,x_k)))"
heat_flux = "Eq(q_j, (1.0/((gama-1)*Minf*Minf*Pr*Re))*Der(T,x_j))"
substitutions = [stress_tensor, heat_flux]
constants = ["Re", "Pr", "gama", "Minf"]

# Formulas for the variables used in the equations
velocity = "Eq(u_i, rhou_i/rho)"
pressure = "Eq(p, (gama-1)*(rhoE - rho*(1/2)*(u_j*u_j)))"
temperature = "Eq(T, p*gama*Minf*Minf/(rho))"
formulas = [velocity, pressure, temperature]

# Create the TGV problem and expand the equations.
problem = Problem(equations, substitutions, ndim, constants=constants,
                  coordinate_symbol="x", metrics=[], formulas=formulas)
expanded_equations = problem.get_expanded(problem.equations)
expanded_formulas = problem.get_expanded(problem.formulas)

spatial_scheme = Central(4) # Fourth-order central differencing in space.
temporal_scheme = RungeKutta(3) # Third-order Runge-Kutta time-stepping scheme.

# Create a numerical grid of solution points
length = [2.0*pi*1.0]*ndim; np = [64]*ndim; deltas = [length[i]/np[i] for i in range(len(length)) ]
grid = Grid(ndim,{ 'delta':deltas, 'number_of_points':np})

# Perform the discretisation
sd = SpatialDiscretisation(expanded_equations, expanded_formulas, grid, spatial_scheme)
td = TemporalDiscretisation(temporal_scheme, grid, constant_dt=True, spatial_discretisation=sd)

# Boundary condition
boundary_condition = PeriodicBoundaryCondition(grid)
for dim in range(ndim):
    boundary_condition.apply(arrays=td.prognostic_variables, boundary_direction=dim)

# Initial conditions
x = "Eq(grid.grid_variable(x), grid.Idx[0]*grid.deltas[0])"
y = "Eq(grid.grid_variable(y), grid.Idx[1]*grid.deltas[1])"
z = "Eq(grid.grid_variable(z), grid.Idx[2]*grid.deltas[2])"
u = "Eq(grid.grid_variable(u),sin(x)*cos(y)*cos(z))"
v = "Eq(grid.grid_variable(v),-cos(x)*sin(y)*cos(z))"
w = "Eq(grid.grid_variable(w), 0.0)"
p = "Eq(grid.grid_variable(p), 1.0/(gama*Minf*Minf)+ (1.0/16.0) * (cos(2.0*x)+cos(2.0*y))*(2.0 + cos(2.0*z)))"
r = "Eq(grid.grid_variable(r), gama*Minf*Minf*p)"
initial_conditions = [x,y,z,u,v,w,p,r, "Eq(grid.work_array(rho), r)", "Eq(grid.work_array(rhou0), r*u)",
                     "Eq(grid.work_array(rhou1), r*v)", "Eq(grid.work_array(rhou2), 0.0)",
                     "Eq(grid.work_array(rhoE), p/(gama-1) + 0.5* r *(u**2+ v**2 + w**2))"]
initial_conditions = GridBasedInitialisation(grid, initial_conditions)

# I/O save conservative variables at the end of simulation
io = FileIO(temporal_discretisation.prognostic_variables)

# Simulation parameters
l1 = ['niter', 'Re', 'Pr', 'gama', 'Minf', 'mu', 'precision', 'name', 'deltat']
l2 = [20000, 1600, 0.71, 1.4, 0.1, 1.0, "double", "taylor_green_vortex", 3.385*10**-3]
simulation_parameters = dict(zip(l1,l2))

# Generate the code.
OPSC(grid, sd, td, boundary_condition, initial_conditions, io, simulation_parameters)

```

Figure A.11: A cut-down version of the 3D Taylor-Green vortex setup/configuration file (67 lines long including whitespace), showing the key components and classes available in OpenSBLI.