Automated derivation and parallel execution of finite difference models on CPUs, GPUs and Intel Xeon Phi processors using code generation

> Christian T. Jacobs, Satya P. Jammy, David J. Lusher, Neil D. Sandham

> > University of Southampton

11 October 2017

Background

- SBLI: CFD code developed at Soton. Written in F95.
- Solves compressible N-S equations:
 - Multi-block
 - 2D & 3D curvilinear grids
 - 4th order central differencing
 - 3rd & 4th order explicit Runge-Kutta timestepping
 - DNS, LES, shock capturing, filtering
 - ~60K lines of code
- Currently capable of running on CPU clusters
 - Scaling to 200,000+ cores (Yao et al., 2009)

Motivation

- Existing archs: multicore CPUs, GPUs, Intel Xeon Phi.
 - Future energy efficient systems from ARM and friends
- Newer hardware has the potential to reduce runtime, but...
- Most models not in a position to readily exploit such architectures to their full potential.
 - Porting SBLI requires non-trivial rewrite.
 - Burden on model devs to not only be domain specialists, but also experts in numerial methods, parallel computing paradigms, and their efficient implementation.
 - Newer architecture might arrive during porting.
 - How do we future-proof codes?

Approach

- Key components of numerical solution are
 - Problem description
 - Numerical method
 - Model code (discretisation, solver, etc) for various architectures



Future-proofing with OPS

- OPS: Oxford Parallel library for Structured-mesh computations
- Multi-block structured applications
- Source-to-source translation for parallel implementations on various architectures
- Very little overhead with the automation process for similar applications (CloverLeaf mini app Mudalige et al. 2014)

Example for simple stencil averaging

ops_par_loop:

int range[4] = {imin,imax,jmin,jmax};

ops_par_loop(calc, block, 2, range,

ops_arg_dat(a,S2D_0,"double",OPS_WRITE),ops_arg_dat(b,S2D_1,"double",OPS_READ));
Kernel:

void calc(double *a, const double *b) {
 a[OPS_ACC0(0,0)] = 0.5*(b[OPS_ACC1(1,0)] + b[OPS_ACC1(-1,0)];)



Source code remains unchanged

Newer architectures backend translator needs to be written

Proof of concept: Shu-Osher case

CPU: Intel[®] Xeon[®] E5-2640 @2.5GHz 12 cores GPU: NVIDIA Tesla K20c 2946 CUDA cores 5GB memory

	CPU		GPU		Speedup
Grid Size (Millions)	12 MPI	12 OpenMP	OpenCL on GPU	CUDA	CPU/GPU
0.0025	2.92	4.61	3.75	3.37	0.89
0.05	33.75	52.89	10.81	9.34	3.61
0.1	66.05	76.47	16.69	16.13	4.09
0.2	136.79	167.13	30.09	29.15	4.69
2	1738.8	2429.44	271.51	264.59	6.57

OPS

- OPS is verbose in nature.
- Error prone.
- No flexibility of numerical method or equations to be solved.



OpenSBLI

- User describes the problem at a higher level.
- Numerical analyst develops the numerical algorithm which generates a sequential model code in OPS-compliant C.
- Computer scientist handles parallel backend implementation.



OpenSBLI: Design

- Written in Python and uses SymPy as building blocks
- User specifies problem in a Python 'problem spec. file'
 - PDEs in Einstein/index notation
 - Spatial scheme and order of accuracy (finite difference)
 - Time-stepping scheme (RK3)
 - Boundary conditions
 - I/O and other parameters for the simulation

OpenSBLI: Design

- Finds and expands the summation indices of PDEs
 - Applies spatial, temporal and boundary schemes
- Creates computational kernels.
- Generates model code. Translation to different arch. with OPS.
- Outputs LaTeX files of the computational kernels for debugging.

OpenSBLI: Design



Example

- 50 line high-level problem definition for the 3D compressible Navier-Stokes equations
- 2000 line generated sequential OPS C code
- 20K lines of generated code for MPI and CUDA

```
# Number of dimensions for the problem ndim = 3
```

```
# Define the compresible Navier-Stokes equations in Einstein notation.
```

```
mass = "Eq(Der(rho,t), - Conservative(rho*u_j,x_j))"
momentum = "Eq(Der(rhou_i,t), -Conservative(rhou_i*u_j + KD(_i,_j)*p ,x_j) + Der(tau_i_j,x_j))"
energy = "Eq(Der(rhoE,t), - Conservative((p+rhoE)*u_j,x_j) + Der(q_j,x_j) + Der(u_i*tau_i_j ,x_j))"
equations = [mass, momentum, energy]
```

```
# Substitutions
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]
```

```
# Define all the constants in the equations
constants = ["Re", "Pr","gama", "Minf"]
```

```
# Define coordinate direction symbol (x) this will be x_i, x_j, x_k
coordinate_symbol = "x"
```

```
# 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]
```

Example

void taylor_green_vortex_block0_69_kernel(const double *wk20 , const double *wk47 , const double *wk21 , const double *wk28 , const double *ul , const double *wk29 , const double *wk19 , const double *wk0 , const double *wk15 , const double *wk35 , const double *wk18 , const double *wk11 , const double *wk12 , const double *wk31 , const double *wk85 , const double *wk37 , const double *wk34 , const double *wk10 , const double *wk30 , const double *wk39 , const double *wk44 , const double *u0 , const double *wk40 , const double *wk46 , const double *wk45 , const double *wk41 , const double *wk25 , const double *wk3 , const double *wk7 , const double *wk46 , const double *wk42 , const double *wk33 , const double *wk6 , const double *wk32 , const double *wk38 , const double *wk14 , const double *wk42 , const double *wk26 , const double *wk43 , const double *wk9 , const double *wk22 , const double *wk27 , const double *wk5 , const double *wk23 , const double *wk9 , const double *wk4 , const double *wk24 , const double *wk27 , const double *wk5 , const double *wk43 , const double *wk9 , const double *wk4 , const double *wk17 , const double *wk13 , const double *wk36 , const double *wk16 , double *wk49 , double *wk48 , double *wk50 , double *wk51 , double *wk52)

```
wk48[0PS ACC52(0,0,0)] = -wk11[0PS ACC11(0,0,0)] - wk14[0PS ACC35(0,0,0)] - wk2[0PS ACC30(0,0,0)];
wk49[0PS ACC51(0,0,0)] = rinv11*(wk0[0PS ACC7(0,0,0)] + wk44[0PS ACC20(0,0,0)]) +
   rinvll*(wk3[OPS ACC27(0,0,0)] + wk47[OPS ACC1(0,0,0)]) + rinvll*((rc4)*wk16[OPS ACC50(0,0,0)] -
   rc6*wk44[0PS ACC20(0,0,0)] - rc6*wk47[0PS ACC1(0,0,0)]) - wk18[0PS ACC10(0,0,0)] - wk20[0PS ACC0(0,0,0)] -
   wk29[0PS ACC5(0,0,0)] - wk39[0PS ACC19(0,0,0)];
wk50[OPS ACC53(0,0,0)] = rinv11*(wk13[OPS ACC48(0,0,0)] + wk42[OPS ACC36(0,0,0)]) +
   rinv11*(wk43[OPS ACC38(0,0,0)] + wk5[OPS ACC43(0,0,0)]) + rinv11*((rc4)*wk26[OPS ACC37(0,0,0)] -
   rc6*wk42[0PS ACC36(0,0,0)] - rc6*wk43[0PS ACC38(0,0,0)]) - wk21[0PS ACC2(0,0,0)] - wk27[0PS ACC42(0,0,0)] -
   wk31[0PS ACC13(0,0,0)] - wk41[0PS ACC25(0,0,0)];
wk51[OPS ACC54(0,0,0)] = rinv11*(wk22[OPS ACC40(0,0,0)] + wk45[OPS ACC24(0,0,0)]) +
   rinvll*(wk46[0PS ACC23(0,0,0)] + wk7[0PS ACC28(0,0,0)]) + rinvll*((rc4)*wk4[0PS ACC46(0,0,0)] -
   rc6*wk45[0PS ACC24(0,0,0)] - rc6*wk46[0PS ACC23(0,0,0)]) - wk28[0PS ACC3(0,0,0)] - wk32[0PS ACC33(0,0,0)] -
   wk36[0PS ACC49(0,0,0)] - wk9[0PS ACC45(0,0,0)];
wk52[0PS ACC55(0,0,0)] = rinvll*rinvl2*rinvl3*rinvl4*wk19[0PS ACC6(0,0,0)] +
   rinv11*rinv12*rinv13*rinv14*wk30[OPS ACC18(0,0,0)] + rinv11*rinv12*rinv13*rinv14*wk35[OPS ACC9(0,0,0)] +
   rinvll*(wk0[0PS ACC7(0,0,0)] + wk44[0PS ACC20(0,0,0)])*u0[0PS ACC21(0,0,0)] +
   rinv11*(wk1[0PS ACC29(0,0,0)] + wk23[0PS ACC44(0,0,0)])*wk1[0PS ACC29(0,0,0)] +
   rinv11*(wk1[0PS ACC29(0,0,0)] + wk23[0PS ACC44(0,0,0)])*wk23[0PS ACC44(0,0,0)] +
   rinv11*(wk12[OPS ACC12(0,0,0)] + wk37[OPS ACC15(0,0,0)])*wk12[OPS ACC12(0,0,0)] +
   rinv11*(wk12[OPS ACC12(0,0,0)] + wk37[OPS ACC15(0,0,0)])*wk37[OPS ACC15(0,0,0)] +
   rinv11*(wk13[OPS ACC48(0,0,0)] + wk42[OPS ACC36(0,0,0)])*u1[OPS ACC4(0,0,0)] +
   rinv11*(wk15[0PS ACC8(0,0,0)] + wk8[0PS ACC14(0,0,0)])*wk15[0PS ACC8(0,0,0)] +
   rinv11*(wk15[OPS ACC8(0,0,0)] + wk8[OPS ACC14(0,0,0)])*wk8[OPS ACC14(0,0,0)] +
   rinv11*(wk22[OPS ACC40(0,0,0)] + wk45[OPS ACC24(0,0,0)])*u2[OPS ACC39(0,0,0)] +
   rinvl1*(wk3[OPS ACC27(0,0,0)] + wk47[OPS ACC1(0,0,0)])*u0[OPS ACC21(0,0,0)] +
   rinv11*(wk43[0PS ACC38(0,0,0)] + wk5[0PS ACC43(0,0,0)])*u1[0PS ACC4(0,0,0)] +
   rinv11*(wk46[0PS ACC23(0,0,0)] + wk7[0PS ACC28(0,0,0)])*u2[0PS ACC39(0,0,0)] +
   rinv11*((rc4)*wk16[OPS ACC50(0,0,0)] - rc6*wk44[OPS ACC20(0,0,0)] -
   rc6*wk47[0PS ACC1(0,0,0)])*u0[0PS ACC21(0,0,0)] + rinv11*(-rc6*wk17[0PS ACC47(0,0,0)] -
   rc6*wk25[0PS ACC26(0,0,0)] + (rc4)*wk34[0PS ACC16(0,0,0)])*wk34[0PS ACC16(0,0,0)] +
   rinvll*(-rc6*wkl7[0PS ACC47(0,0,0)] + (rc4)*wk25[0PS ACC26(0,0,0)] -
   rc6*wk34[0PS ACC16(0,0,0)])*wk25[0PS ACC26(0,0,0)] + rinv11*((rc4)*wk17[0PS ACC47(0,0,0)] -
   rc6*wk25[0PS ACC26(0,0,0)] - rc6*wk34[0PS ACC16(0,0,0)])*wk17[0PS ACC47(0,0,0)] +
   rinv11*((rc4)*wk26[0PS ACC37(0,0,0)] - rc6*wk42[0PS ACC36(0,0,0)] -
   rc6*wk43[0PS ACC38(0,0,0)])*u1[0PS ACC4(0,0,0)] + rinv11*((rc4)*wk4[0PS ACC46(0,0,0)] -
   rc6*wk45[0PS ACC24(0,0,0)] - rc6*wk46[0PS ACC23(0,0,0)])*u2[0PS ACC39(0,0,0)] - wk10[0PS ACC17(0,0,0)] -
   wk24[OPS ACC41(0,0,0)] - wk33[OPS ACC31(0,0,0)] - wk38[OPS ACC34(0,0,0)] - wk40[OPS ACC22(0,0,0)] -
   wk6[0PS ACC32(0,0,0)];
```

Example of auto-generated kernel for computing RHS of compressible Navier-Stokes equation.

OpenSBLI: Advantages

- Flexible choice of equations
- Spatial order can be easily varied
- Implementing new numerical method requires symbolic representation
 - To generate a Fortran or code in another language, only a new OPS backend need be written.

Verification & Validation





Solution convergence for the 2D advection-diffusion equation (method of manufactured solutions). Image by Jacobs et al. (2017). Used under CC-BY licence.

Verification & Validation

- 3D Taylor-Green vortex problem N-S Equations
- Up to 1 billion grid points,
- Re = 1600
- CPU(ARCHER), GPU(K40c)

Images by Jacobs et al. (2017). Used under CC-BY licence.





Flexibility of algorithms

- Future HPC machines expected to deliver exascale capabilities
- Theoretical flops of many-core processors (e.g. GPUs, Xeon Phi cards) have increased but RAM is limited.
- Requires novel algorithmic changes to exploit the flops
- Reducing memory usage also reduces the amount of data to be transferred between CPU and GPU.
- A **detailed study** of the algorithm's performance in the context of compressible Navier-Stokes equations is important because of the uncertainty of these architectures.

Flexibility of algorithms

Using the OpenSBLI framework we can compare different algorithmic choices without need to rewrite the code

- **Baseline (BL)** all the derivatives are stored in work arrays
- Recompute All (RA) all continuous derivatives in the governing equations are replaced by their discretised formula
- Recompute Some (RS) only the first derivatives of velocity are stored and the rest recomputed
- Store None (SN) all the derivatives are evaluated as thread/process local variables
- Store Some (SS) only the first derivatives of velocity are stored and the rest are evaluated as thread/process local variables

Algorithmic performance

	ARCHER node (24 MPI processes), runtime (s)				
Grid Size (Millions)	Baseline	Recompute All	Recompute Some	Store None	Store Some
0.2	16	9	11	8	10
2.09	183	98	97	91	89
16.77	1562	765	803	694	685

	CUDA Tesla K40c, runtime (s)					
Grid Size (Millions)	Baseline	Recompute All	Recompute Some	Store None	Store Some	
0.2	9	6	6	6	5	
2.09	57	35	35	41	33	
16.77	495	259	256	302	246	

Parallel scaling



Images by Jammy et al. (In Press). Used under CC-BY licence.

CPU to GPU speed-up



CPU to GPU speed up 2.7~3.15

Images by Jammy et al. (In Press). Used under CC-BY licence.

Intel Xeon Phi (KNL) performance

	ARCHER node (24 MPI processes), runtime (s)				
Grid Size (Millions)	Baseline	Recompute All	Recompute Some	Store None	Store Some
0.2	16	9	11	8	10
2.09	183	98	97	91	89
16.77	1562	765	803	694	685

	Intel KNL (64 MPI processes), runtime (s)				
Grid Size (Millions)	Baseline	Recompute All	Recompute Some	Store None	Store Some
0.2	20	13	14	13	13
2.09	148	92	107	87	102
16.77	688	675	647	610	562

Advantages and limitations

- New DSLs can be readily integrated in OpenSBLI to aid futureproofing
- Flexibility of algorithms, methods and equations.
- Use of external libraries like FFTW & implementation of implicit solvers needs to be done at OpenSBLI and the backend OPS level.
- Debugging for errors at different levels (partly helped by the OpenSBLI LaTeX debugging).

Conclusions & Acknowledgements

- New framework for the automated solution of finite difference methods on various architectures is developed and validated.
- Separation of concerns enables better model maintainability, and future-proofs the code as newer architectures arrive.
- Storing the first derivatives of velocity in the context of compressible Navier-Stokes solution is optimal across architectures.
- Funded by EPSRC grants EP/K038567/1 and EP/L000261/1, and European Commission H2020 grant 671571 "ExaFLOW: Enabling Exascale Fluid Dynamics Simulations"

Resources

- OpenSBLI released under GNU GPL: https://opensbli.github.io
- C. T. Jacobs, S. P. Jammy, N. D. Sandham (2017). OpenSBLI: A framework for the automated derivation and parallel execution of finite difference solvers on a range of computer architectures. *Journal of Computational Science*. DOI: 10.1016/j.jocs.2016.11.001
- S. P. Jammy, C. T. Jacobs, N. D. Sandham (In Press). Performance evaluation of explicit finite difference algorithms with varying amounts of computational and memory intensity. *Journal of Computational Science*. DOI: 10.1016/j.jocs.2016.10.015