

# PyFR: Heterogeneous Computing from a Homogeneous Codebase

P. E. Vincent

Department of Aeronautics  
Imperial College London



20<sup>th</sup> November 2014 (Antony's Birthday!)



# Overview

- Motivation
- Flux Reconstruction
- Modern Hardware
- PyFR
- Results
- Summary



# Motivation



Current industry standard CFD tools  
have limited capabilities



# Motivation





Technology is decades old and designed for solving steady flow problems (using RANS approach)





# Motivation





Technology is decades old and designed for solving steady flow problems (using RANS approach)



# Motivation



Need to expand the 'industrial CFD envelope'



# Motivation

- “reliable use of CFD has remained confined to a small but important region of the operating design space due to the inability of current methods to reliably predict turbulent separated flows” [2]

# Motivation

- Objective of our research is to advance industrial CFD capabilities from their current 'RANS plateau'

# Motivation





- We aim to develop the *de facto* industry standard technology for **affordable** (and hence industrially relevant) high-fidelity scale-resolving simulations of **unsteady flow phenomena** within the vicinity of **complex geometric configurations**

# Motivation

- Achieved by intelligently leveraging benefits of (and synergies between) high-order **Flux Reconstruction (FR)** methods for unstructured grids and massively-parallel **modern hardware platforms**

# Motivation

Flux Reconstruction  
+  
Modern Hardware



# Flux Reconstruction

- Flux Reconstruction (FR) approach to high-order methods was first proposed by Huynh in 2007 [3]
- High-order accurate in space
- Works on unstructured grids

# Flux Reconstruction

- So ...

High Accuracy + Complex Geometry

# Flux Reconstruction

- Nature of FR scheme depends on location of solution points, interface flux, correction function
- Can recover a wide range of schemes via judicious choice of correction function [4]
- A one-parameter family of provably stable FR schemes have been identified [5]

[4] H.T. Huynh. A flux Reconstruction Approach to High-Order Schemes Including Discontinuous Galerkin Methods. AIAA Paper 2007-4079. 2007


[5] P. E. Vincent, P. Castonguay, A. Jameson. A New Class of High-Order Energy Stable Flux Reconstruction Schemes. Journal of Scientific Computing. 2011

# Flux Reconstruction

- We have recently identified a new multi-parameter family of stable FR schemes


# Flux Reconstruction

- Plot **stable** region of parameter space ...




# Flux Reconstruction

- Original one-parameter family ...




# Flux Reconstruction

- Perform a whole bunch of numerical experiments ...



# Flux Reconstruction

- And record which ones were **stable/unstable** ...




# Modern Hardware



# Modern Hardware

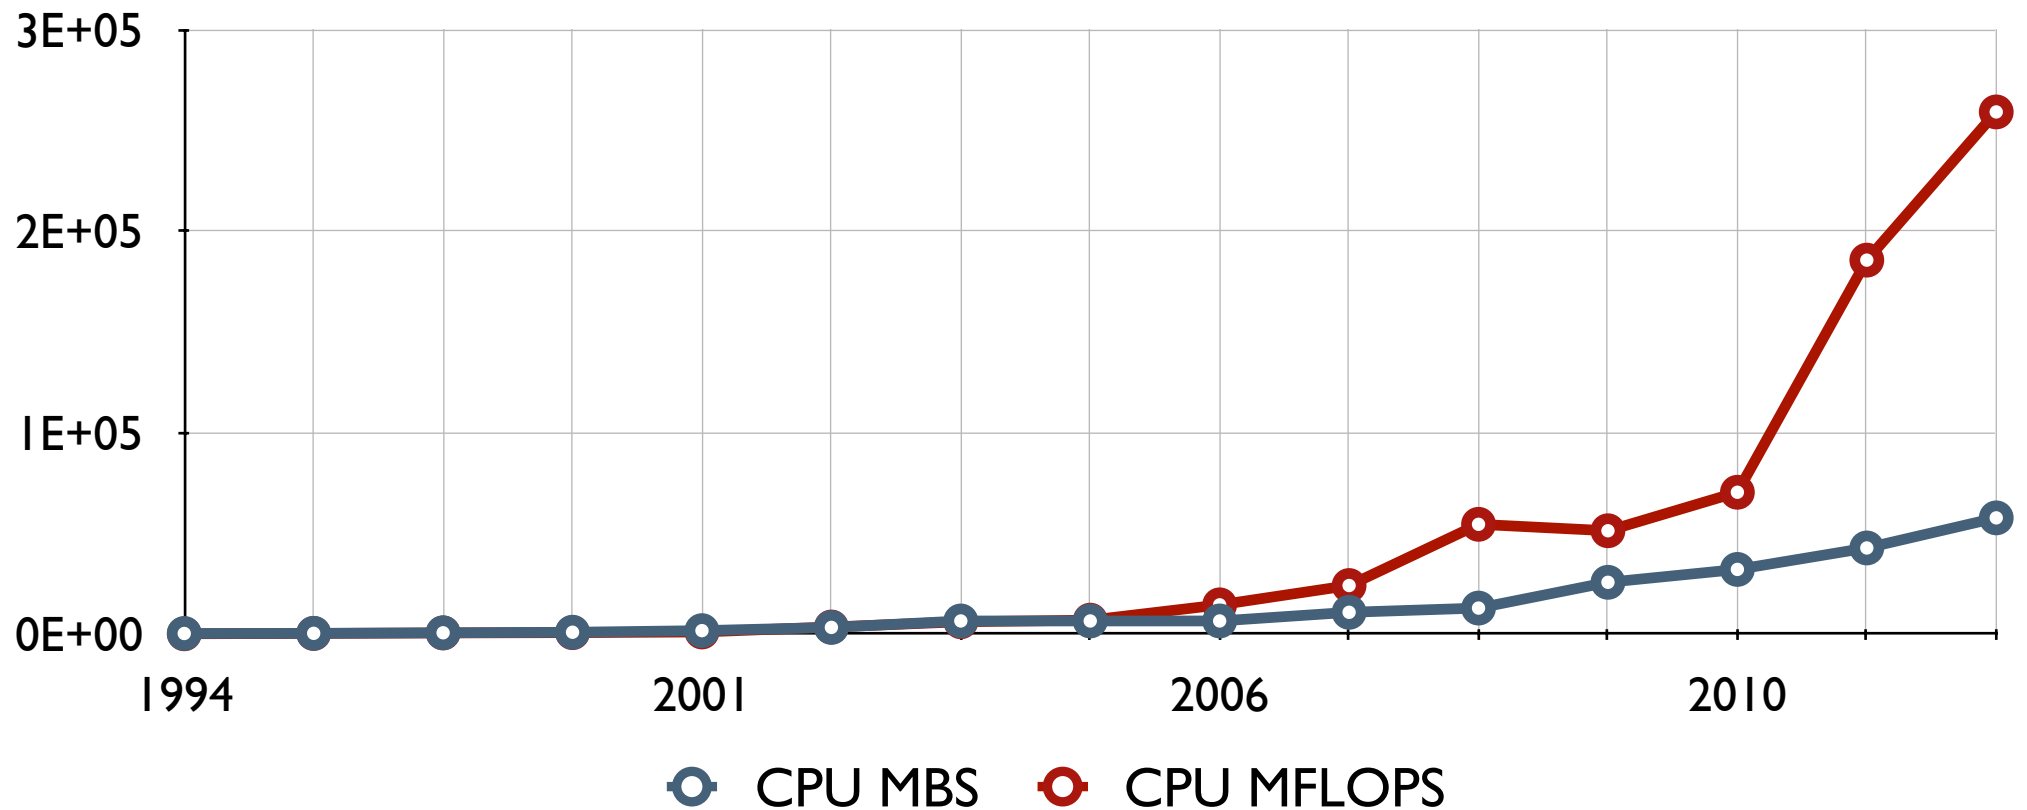


# Modern Hardware



# Modern Hardware

- FLOPS increasing faster than memory bandwidth



[2] F. D. Witherden, A. M. Farrington, P. E. Vincent. PyFR: An Open Source Framework for Solving Advection-Diffusion Type Problems on Streaming Architectures using the Flux Reconstruction Approach. Computer Physics Communications. 2014. Data courtesy of Jan Treibig.

# Modern Hardware

- Also FLOPS come in parallel ...

# Modern Hardware



- Multiple **Cores** each with **Vector Unit**
- Parallelism exposed via **MIMD** (Multiple Instruction Multiple Data) + **SMT** (Simultaneous Multi Threading) + **SIMD** (Single Instruction Multiple Data) + **ILP** (Instruction Level Parallelism)

# Modern Hardware



- Nvidia GPUs - multiple Streaming Multi-Processors each with CUDA Cores
- Parallelism exposed via SIMT (Single Instruction Multiple Thread)

# Modern Hardware



- AMD GPUs - multiple Compute Units each with Stream Processors
- Parallelism exposed via SIMT (Single Instruction Multiple Thread)

# Modern Hardware


- Also, different programming languages for different devices

# Modern Hardware

- So a **challenging** environment ...

# Modern Hardware

- So a **challenging** environment ...



# Modern Hardware




- But significant **FLOPS** now available if they can be harnessed ...

2.6TFLOPS  
(Double Precision)






# PyFR

Flux Reconstruction  
+  
Modern Hardware



# PyFR

PyFR



# PyFR

- Features (v0.2.3 - current release)

<b>Governing Equations</b>	Compressible Euler Compressible Navier Stokes
<b>Spatial Discretisation</b>	Arbitrary order FR on mixed unstructured grids (tris, quads, hexes, tets, prisms)
<b>Temporal Discretisation</b>	Range of explicit Runge-Kutta schemes
<b>Platforms</b>	CPU clusters (C-OpenMP-MPI) Nvidia GPU clusters (CUDA-MPI) AMD GPU clusters (OpenCL-MPI)
<b>Precision</b>	Single Double
<b>Input</b>	Gmsh
<b>Output</b>	Paraview

# PyFR

- Features (v0.2.3 - current release)

Governing Equations	Compressible Euler Compressible Navier Stokes
Spatial Discretisation	Arbitrary order FR on mixed unstructured grids (tris, quads, hexes, tets, prisms)
Temporal Discretisation	Range of explicit Runge-Kutta schemes
Platforms	CPU clusters (C-OpenMP-MPI) Nvidia GPU clusters (CUDA-MPI) AMD GPU clusters (OpenCL-MPI)
Precision	Single Double
Input	Gmsh
Output	Paraview

# PyFR

- Python Outer Layer (Hardware Independent)

Python Outer Layer  
(Hardware Independent)

- Setup
- Distributed memory parallelism
- Outer 'for' loop and calls to Hardware Specific Kernels

# PyFR

- Need to generate the Hardware Specific Kernels

## Python Outer Layer (Hardware Independent)

- Setup
- Distributed memory parallelism
- Outer 'for' loop and calls to Hardware Specific Kernels

# PyFR

- Two **types** of kernel are required ...

## Python Outer Layer (Hardware Independent)

- Setup
- Distributed memory parallelism
- Outer 'for' loop and calls to **Hardware Specific Kernels**

## Matrix Multiply Kernels

- Data interpolation/  
extrapolation  
etc.

## Point-Wise Nonlinear Kernels

- Flux functions,  
Riemann  
solvers etc.

# PyFR

- For matrix multiply kernels it is pretty easy ...

## Python Outer Layer (Hardware Independent)

- Setup
- Distributed memory parallelism
- Outer 'for' loop and calls to Hardware Specific Kernels

## Matrix Multiply Kernels

- Data interpolation/  
extrapolation  
etc.

## Point-Wise Nonlinear Kernels

- Flux functions,  
Riemann  
solvers etc.

Use DGEMM from  
vendor supplied  
BLAS

# PyFR

- Harder for point-wise nonlinear kernels ...

## Python Outer Layer (Hardware Independent)

- Setup
- Distributed memory parallelism
- Outer 'for' loop and calls to Hardware Specific Kernels

## Matrix Multiply Kernels

- Data interpolation/  
extrapolation  
etc.

## Point-Wise Nonlinear Kernels

- Flux functions,  
Riemann  
solvers etc.

Use DGEMM from  
vendor supplied  
BLAS

Pass Mako  
derived kernel  
templates through  
Mako derived  
templating engine

# PyFR

- These can now be called

## Python Outer Layer (Hardware Independent)

- Setup
- Distributed memory parallelism
- Outer 'for' loop and calls to Hardware Specific Kernels

## Matrix Multiply Kernels

- Data interpolation/  
extrapolation  
etc.

## Point-Wise Nonlinear Kernels

- Flux functions,  
Riemann  
solvers etc.

## C/OpenMP Hardware Specific Kernels



## CUDA Hardware Specific Kernels



## OpenCL Hardware Specific Kernels



Use DGEMM from  
vendor supplied  
BLAS

Pass Mako  
derived kernel  
templates through  
Mako derived  
templating engine

# PyFR

- These can now be called

## Python Outer Layer (Hardware Independent)

- Setup
- Distributed memory parallelism
- Outer 'for' loop and calls to Hardware Specific Kernels

## Matrix Multiply Kernels

- Data interpolation/  
extrapolation  
etc.

## Point-Wise Nonlinear Kernels

- Flux functions,  
Riemann  
solvers etc.

C/OpenMP  
Hardware  
Specific  
Kernels



CUDA  
Hardware  
Specific  
Kernels



OpenCL  
Hardware  
Specific  
Kernels



Use DGEMM from  
vendor supplied  
BLAS

Pass Mako  
derived kernel  
templates through  
Mako derived  
templating engine

# PyFR

- ~5.5k lines of Python

# PyFR

- Open source '3 Clause New Style BSD License'

# PyFR



- Website: [www.pyfr.org](http://www.pyfr.org)
- Twitter: [@PyFR\\_Solver](https://twitter.com/PyFR_Solver)
- Paper: [Computer Physics Communications \[3\]](#)

[3] F. D. Witherden, A. M. Farrington, P. E. Vincent. PyFR: An Open Source Framework for Solving Advection-Diffusion Type Problems on Streaming Architectures using the Flux Reconstruction Approach. Accepted for publication in Computer Physics Communications. 2014


# Results

- Taylor-Green vortex breakdown
- $Re = 1600$
- $Ma = 0.1$
- Compare with van Rees et al. [7]

[7] W. M. van Rees, A. Leonard, D. I. Pullin, and P. Koumoutsakos. A Comparison of Vortex and Pseudo-Spectral Methods for the Simulation of Periodic Vortical Flows at High Reynolds Numbers. *Journal of Computational Physics*, 2011


# Results

- A movie ...




# Results

- A movie ...




# Results

- van Rees et al. spectral DNS + PyFR (2nd order hex)




# Results

- van Rees et al. spectral DNS + PyFR (3rd order hex)




# Results

- van Rees et al. spectral DNS + PyFR (4th order hex)




# Results

- van Rees et al. spectral DNS + PyFR (5th order hex)




# Results

- van Rees et al. spectral DNS + PyFR (6th order hex)



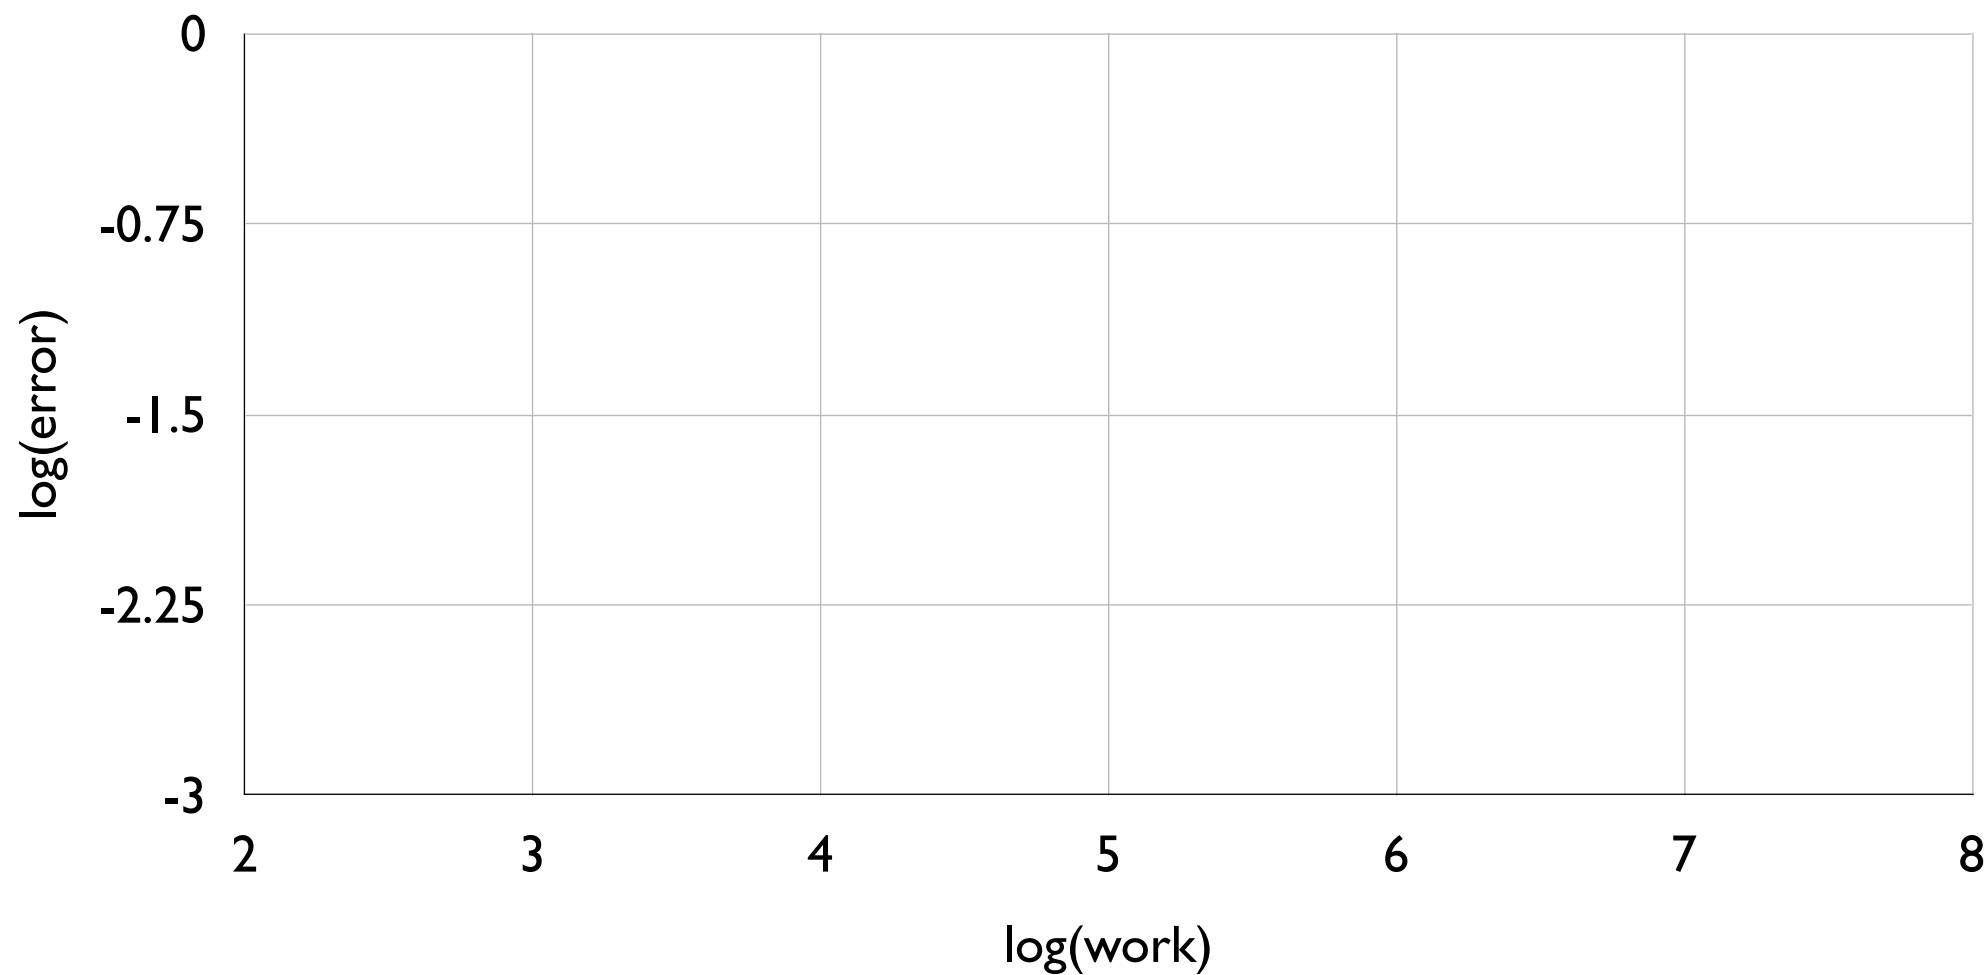
# Results

- van Rees et al. spectral DNS + PyFR (6th order hex)




# Results

- $L_\infty$  error in decay rate




# Results

- $L_\infty$  error in decay rate




# Results

- $L_\infty$  error in decay rate




# Results

- $L_\infty$  error in decay rate




# Results

- $L_\infty$  error in decay rate




# Results

- $L_\infty$  error in decay rate




# Results

- $L_\infty$  difference between decay rate and enstrophy




# Results

- $L_\infty$  difference between decay rate and enstrophy




# Results

- $L_\infty$  difference between decay rate and enstrophy




# Results

- $L_\infty$  difference between decay rate and enstrophy




# Results

- $L_\infty$  difference between decay rate and enstrophy



# Results

- $L_\infty$  difference between decay rate and enstrophy




# Results

- Flow over a **circular cylinder**
- $Re = 3900$
- $Ma = 0.2$
- Compare with Parnaudeau et al. [6]


# Results

- A movie ...




# Results

- A movie ...




# Results

- Parnaudeau et al. experiment




# Results

- Parnaudeau et al. experiment + Parnaudeau et al. LES




# Results

- Parnaudeau et al. experiment + PyFR (5th order hex)




# Results

- Parnaudeau et al. experiment




# Results

- Parnaudeau et al. experiment + Parnaudeau et al. LES




# Results

- Parnaudeau et al. experiment + PyFR (5th order hex)




# Results

- Parnaudeau et al. experiment




# Results

- Parnaudeau et al. experiment + Parnaudeau et al. LES



# Results

- Parnaudeau et al. experiment + PyFR (5th order hex)




# Results

- Flow over a **NACA 0021** at 60 degree AoA
- $Re = 270,000$
- $Ma = 0.2$


# Results

- A movie ...




# Results

- A movie ...




# Results

- Performance on a heterogeneous workstation




# Results

- Single-node on prism/tetrahedral mesh



# Results

- Multi-node heterogeneous on prism/tetrahedral mesh







# Summary

Flux Reconstruction  
+  
Modern Hardware



# Summary

PyFR



# Team



Brian Vermeire



Antony Farrington



Lorenza Grechy



Freddie Witherden



George Ntemos




Francesco Iori



Jin Seok Park

# Funding



# Project Partners



# Emerald



- The authors would like to acknowledge the work presented here made use of the **EMERALD** HPC facility provided by the Centre for Innovation

# Antony

$x \cdot x \cdot \dots \cdot x \cdot v \cdot f$

$$SD \quad \frac{du}{dt} + D_{jk} S_k + \hat{L}'(x) (\hat{F} - f) = 0 \quad f = av \quad v = \sum_{j=1}^n v_j \hat{L}_j \quad f = \sum_{j=1}^n f_j \hat{L}_j \quad n = p+1$$

$S_1 = \hat{F}_1, f_{n+1} = \hat{F}_{n+1}$

Multiply by the standard mass matrix  $M_{jk} = \int \hat{L}_j \hat{L}_k dx \Rightarrow MD = S$

$$\sum_k M_{jk} \frac{du_k}{dt} + \sum_k S_{jk} S_k + \sum_k M_{jk} \hat{L}'(x_k) (\hat{F} - f) = 0.$$

$$\sum_k M_{jk} \hat{L}'(x_k) = \int \hat{L}_j \left( \sum_k \hat{L}_k \hat{L}'(x_k) \right) dx = \int \hat{L}_j \hat{L}'_1 dx = \hat{L}_j \hat{L}'_1 = \hat{L}'_1 \hat{L}_j = -\hat{L}'_1(x) - \hat{L}'_1 \hat{L}_j dx$$

$$\sum_k M_{jk} \frac{du_k}{dt} + \sum_k S_{jk} S_k + \hat{L}'_1 (\hat{F} - f) = 0$$

With upwind flux last term is  $-\hat{L}'_1 (\hat{F} - f)$  The difference is  $-\int \hat{L}_j \hat{L}'_1 dx (\hat{F} - f)$

If we set  $\hat{L}_1 = (-1)^p \frac{1}{2} (1-x) L_p(x)$  by choosing flux points as zeros of  $L_p$ ,

$$\int \hat{L}_1 \hat{L}'_1 dx = (-1)^p \int \frac{1}{2} (1-x) L_p' dx = (-1)^{p+1} \int L_p x dx$$

by orthogonality because  $L_p$  is orthogonal.

We need to choose  $Q = M + c dt$  where  $d$  is  $p$ th diff operator

so  $d^T D f = 0$  for polynomials of degree  $p$   $(M + c dt) D = S$  so that

$$c dt \hat{L}'_1 = c dt \hat{L}_1^{(p+1)} = (-1)^p \int L_p x dx$$

$a_p dt + \dots$  then  $d_p = \frac{L_p^{(p+1)}}{p!} = p! a_p \quad \int L_p x dx = \frac{2^{p+1}}{2^{p+1} p!} \frac{1}{p+1} p^{p+1}$

$L_p = a_p x^p + \dots$ , then  $\hat{L}_1^{(p+1)} = (p+1)! c_p, (c dt \hat{L}_1^{(p+1)})_1 = c p! a_p (p+1)! c_p$

Now is  $\frac{1}{2} \int (v^2 + c u^{(p)}) dx$

$c = \frac{1}{2^p} \frac{p! (p+1)!}{p! (p+1)! 2^{p+1}}$

