A Robust Technique to Make a 2D Advection Solver Tolerant to Soft Faults

Peter Strazdins [2] Brendan Harding [1] Brian Lee [2]
Jackson Mayo [3] Jaideep Ray [3] Robert C. Armstrong [3]

[1]Mathematical Sciences Institute, Australian National University
Canberra, ACT, Australia

[2]Research School of Computer Science, Australian National University
Canberra, ACT, Australia

[3]Sandia National Laboratories
Livermore, California, USA

ICCS - 7th June 2016

Outline

Faults in High Performance Computing

- A challenge for HPC is faults
- Exposure/risk increases with system size
- Many types: hard, soft, network, silent
- Many causes: hardware, software, radiation, network, etc.
- generic solutions: triple modular redundancy (TMR), checkpoint-restart

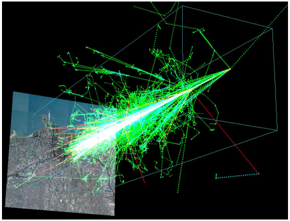

(source: wikipedia)

- Active research area in recent decades
- Numerous papers discuss use of checksums to detect and correct memory failures (bit flips) in linear algebra
- We present an approach for avoiding bit flip errors in finite difference computations

Finite difference computations

- Finite differences methods and method of lines are common for solving partial differential equations
- Explicit methods cannot leverage fault tolerant linear algebra

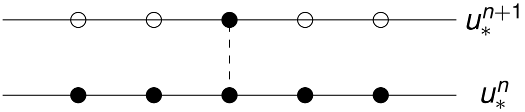

- Triple modular redundancy (TMR) could easily be used
  - Do everything 3 times (with separate memory)
  - Choose result which is equal for any two
    - $1/3$ efficiency (3 times the resources/time)
- How else could we detect/correct or even avoid errors?

Robust stencils in 1D

- Consider the advection equation

$$\partial_t u + a\,\partial_x u = 0$$

- The standard Lax–Wendroff method

$$u_i^{n+1} = \frac{c(1+c)}{2} u_{i-1}^n + (1-c^2) u_i^n + \frac{c(-1+c)}{2} u_{i+1}^n$$

  with $c = a\Delta t/\Delta x$ is stable and second order.
- Ray, Mayo and Armstrong considered using several finite difference discretisations having distinct stencils to make computations fault tolerant.

- One introduces the widened discretisation

$$u_i^{n+1} = \frac{c(2+c)}{8} u_{i-2}^n + \frac{4-c^2}{4} u_i^n + \frac{c(-2+c)}{8} u_{i+2}^n$$

  and a third (far) discretisation

$$u_i^{n+1} = \frac{-3+8c+3c^2}{48} u_{i-3}^n + \frac{9-c^2}{16} u_{i-1}^n + \frac{9-c^2}{16} u_{i+1}^n + \frac{-3-8c+3c^2}{48} u_{i+3}^n$$

  both of which are also stable and second order.
- The corresponding stencils are:

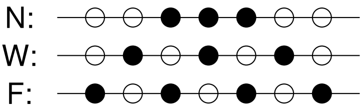

Error avoidance and results in 1D

- To avoid (significant) errors:
  - take the median of $u^n_{i-1}, u^n_i, u^n_{i+1}$
  - discard the $u^n_{i-3}, u^n_{i-2}, \ldots, u^n_{i+3}$ furthest from median
  - Use most compact stencil which avoids discarded value

Results:

- Similar robustness to TMR
- Similar efficiency to TMR (note: not yet optimised)
- Similar ideas also applied to (inviscid) Burgers' equation with similar success for robustness (note: shocks still captured)

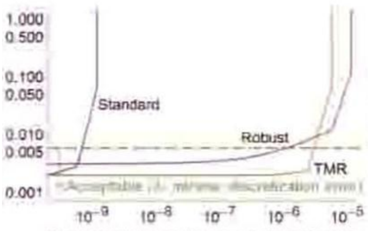

Advection in multiple dimensions

- We extend this work to 2D

$$\partial_t u + a\,\partial_x u + b\,\partial_y u = 0$$

  (NB: also applies to $d > 2$)
- Consider a square domain discretised as a uniform grid
- The 1D discretisations can be applied in 2D by applying along one direction at a time (effectively an operator splitting approach)
- Resulting stencils are on the right

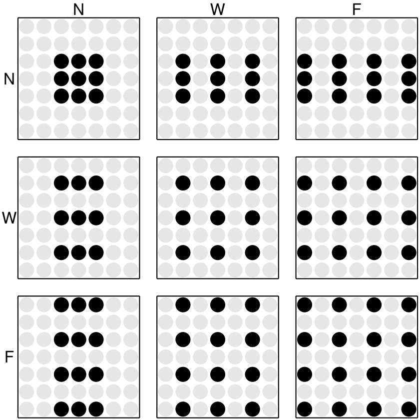

Robust stencils in multiple dimensions

- How do we choose one of these stencils in a way which avoids (significant) faults?
- Computing a median of central values and eliminating outliers from the 7x7 stencil region is expensive
- Instead we consider computing a subset of the 9 stencils and taking the median update as the result
- We consider subsets of 3, 5 and 7 stencils (to simplify the calculation of a median)
- By choosing $2n + 1$ stencils such that no one neighbour is contained in more than n of the stencils then the medain update will avoid any significant faults

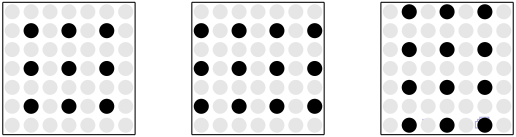

Examples of robust stencil sets

- Ideally we want the NN stencil in our set (most accurate)
- Additionally, the other stencils will bracket this in smooth error free regions, in practice this is tricky
- As a heuristic we choose symmetric stencil sets (i.e. if $XY \in S$ then $YX \in S$ where $X, Y \in \{N, W, F\}$)
- There's one such set (having 5 stencils)
- 3 and 7 stencil sets (not using NN) are also depicted

	N	W	F
N			
W		*	*
F		*	

	N	W	F
N	*		
W		*	*
F		*	*

	N	W	F
N		*	*
W	*	*	*
F	*	*	

Rough performance analysis of methods

Comparisons relative to standard Lax-Wendroff

TMR

- 3x memory
- 3x computation (plus agreement/median)
- 3x boundary comm.

Robust Stencils (3/5/7 examples)

- $\approx$1x memory (extra ghost/halo points)
- 3.7/6.4/8.3x computation (plus median)
- $\approx$3x boundary comm. (wider halos)

- Finite difference methods are often memory bound
- Extra computations for robust stencils potentially significantly cheaper with optimised cache usage

Rough analysis of robustness

Both approaches survive single fault in any one time step

Suppose 2 faults have occurred in a single time step:

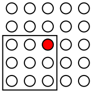

- TMR fails if 2nd fault within 5x5 region around 1st (on another copy, thus 2 chances)
- RS fails if 2nd fault within 13x13 region around 1st (worst case)
- Assuming fault rate $\propto$ memory in use then TMR has 3 times the exposure/risk
- Thus approximately 150 vs 168 elements of exposure/risk of unavoidable failure

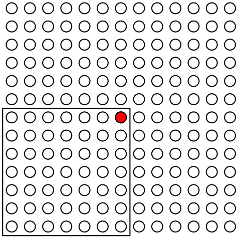

How do we inject faults

We use an additional process to simulate faults

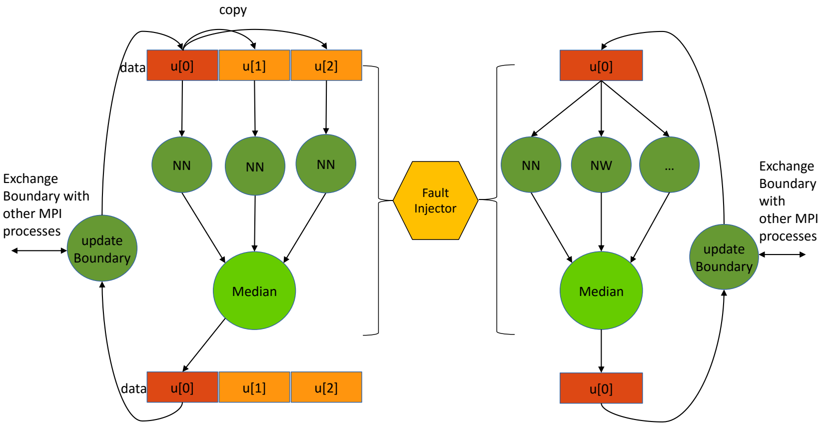

Time between faults is assumed to be exponentially distributed

Other numerical implementation details

- Code is MPI parallel and scalable to at least 2K proc.
- Tests performed on Raijin cluster (NCI, Canberra, Aus.) with gnu compilers
- Codes not yet optimised (beyond twiddling compiler flags)
- Initial condition is a sinusoidal field which is translated periodically via advection
- Faults are injected from another thread at a specified rate, location of faults in memory is uniformly random

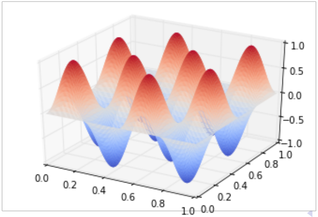

Numerical results

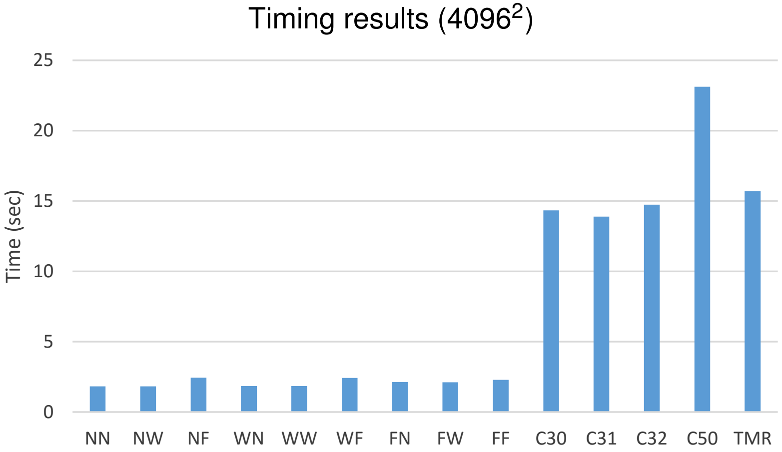

Each takes similar time on own, 3 sets slightly better than TMR

Numerical results

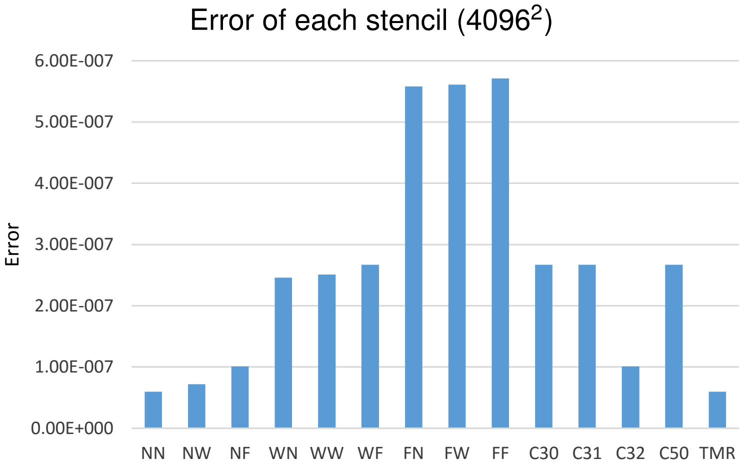

Robust stencil results typically similar to the WW stencil

Serial fault injection results (512^2)

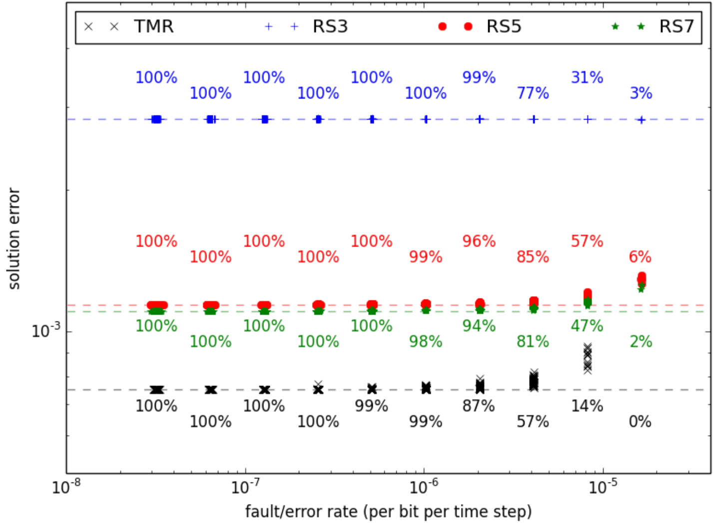

Summary

- Robust stencils are competitive with TMR - similar robustness at a lower computational cost
- Can be extended to any number of dimensions
- Can be extended to other finite difference discretisations
- Works in a select case of finite element methods (e.g. by recasting as a finite difference method)
- Main idea could be applied more generally to other methods
- A detailed analysis is ongoing work
- With multistep integrators one can potentially avoid errors that effect large regions of bits

Acknowledgements

This work was partially supported under Australian Research Council's Linkage Grant scheme (project LP110200410)

This research was undertaken with the assistance of resources from the National Computational Infrastructure (NCI), which is supported by the Australian Government

Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000