Towards real-time simulations of Cardiac Arrhythmias

Lawrence Mitchell
EPCC
lawrence.mitchell@ed.ac.uk

Overview

- CARP is a finite element code that models cardiac electrophysiology
- In-silico modelling of treatments for cardiovascular disease
- Cardiac resynchronisation therapy
- Drug tests
- Defibrillation

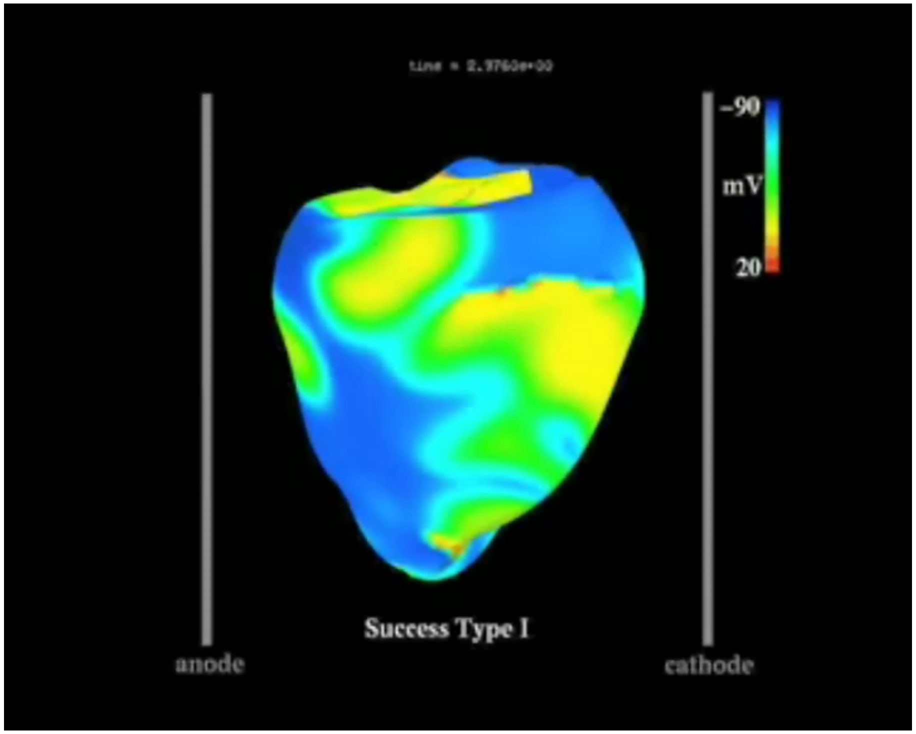

J. Eason and C. Glisson (scholarpedia.org/article/Bidomain_model)

How does it work?

- MRI data converted into unstructured finite element mesh
- Solve set of coupled PDEs and ODEs on this mesh.

Elliptic PDE:

$$(A_i + A_e)\Phi_e^{k+1} = A_i V^{k+1} + I_e$$

Parabolic PDE:

$$\begin{cases} V^{k*} = (1 - \Delta t A_i)\, V^k - \Delta t A_e \phi_e^k & \Delta x > 100\mu m \\ \left[1 + \frac{1}{2}\Delta t A_i\right] V^{k*} = \left[1 - \frac{1}{2}\Delta t A_i\right] V^k - \Delta t A_e \phi_e^k & \Delta x < 100\mu m \end{cases}$$

ODE's:

$$\begin{aligned} V^{k+1} &= V^{k*} + \frac{\Delta t}{C_m} i_{ion}\left(V^{k*}, \vec{\eta}^k\right) \\ \vec{\eta}^{k+1} &= \vec{\eta}^k + \Delta t\, g(V^{k+1}, \vec{\eta}^k) \end{aligned}$$

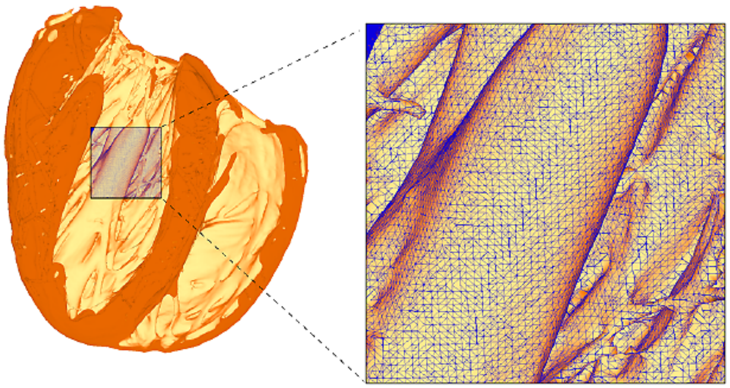

Figure courtesy G. Plank

Performance

- To match experimental results requires high resolution models, upwards of 10 million unknowns.
- Previous best efforts show reasonable parallel efficiency only to 64 cores (~7 million unknowns)

N_p #	Parabolic [s]	ODE [s]	Elliptic [s]	Total Mono [s]	Total Bidomain [s]
4	73.03	7.46	X^*	80.49	X^*
8	42.50	3.72	401.82	46.22	448.04
16	21.84	1.86	214.34	23.70	238.04
32	13.23	0.94	114.33	14.17	128.50
64	6.32	0.39	63.89	6.70	70.59
128	4.03	0.17	61.01	4.22	65.23

Plank *et al*, Phil Trans A (2009)

Poor load balance

- PDEs require large sparse matrix-vector multiply.
- ODEs are local
- Need to minimise partition interfaces constrained by equal-sized partitions
- Default CARP partitioning does neither of these things

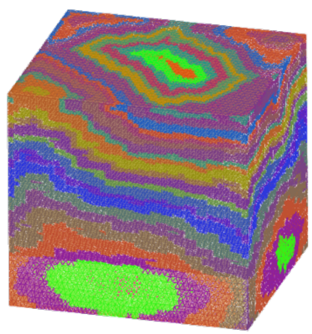

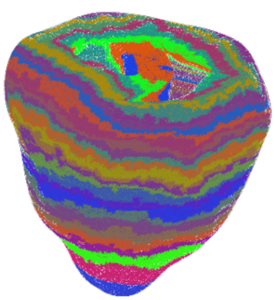

Repartitioning

- Mesh is static, so just repartition prior to constructing model
- We used ParMETIS
- Much better load balance

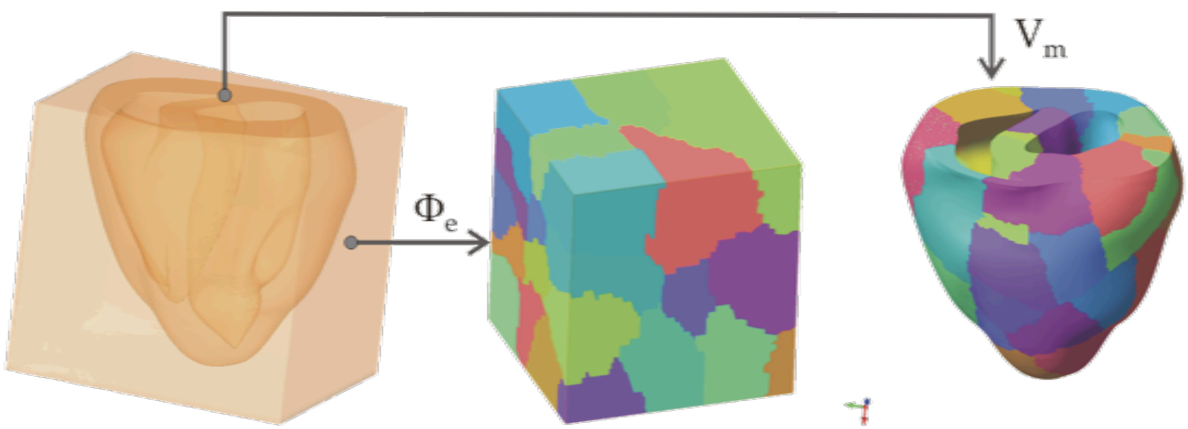

Big win

- Better strong scaling, and better performance
- Rabbit heart: 6.9 million unknowns, 40 million elements

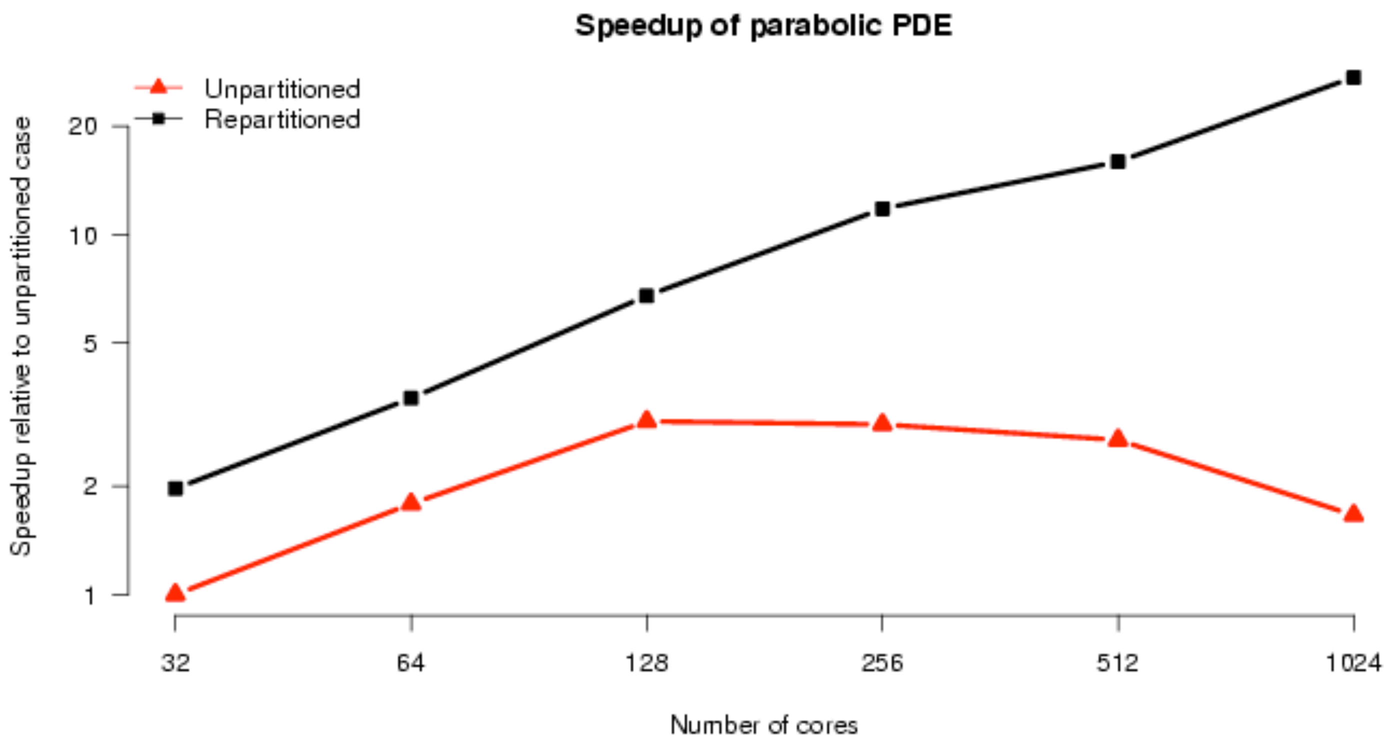

- Elliptic PDE does less well
- Algebraic multigrid preconditioner doesn't scale well enough. Coarse grid matrices are too dense.

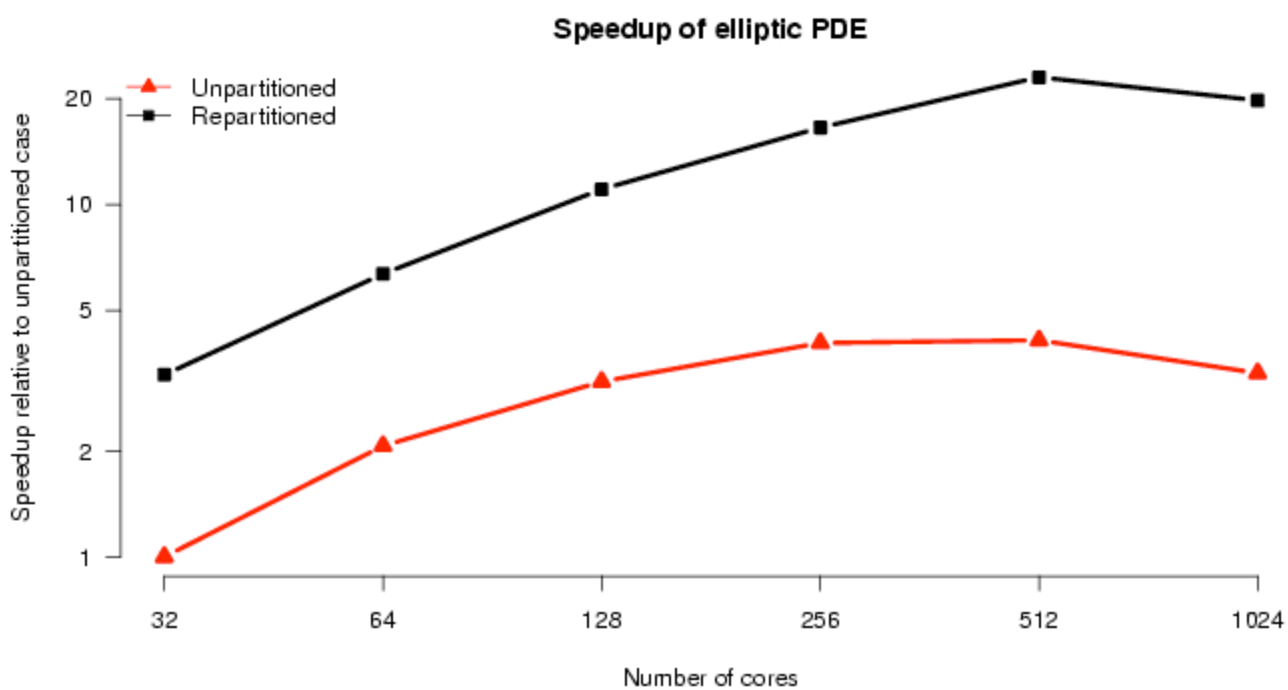

Output hurts

- Profile to find hotspots. Output is serialised, bad scaling.
- Additional problem: partitioning adds extra output latency --- need to map data back onto canonical mesh

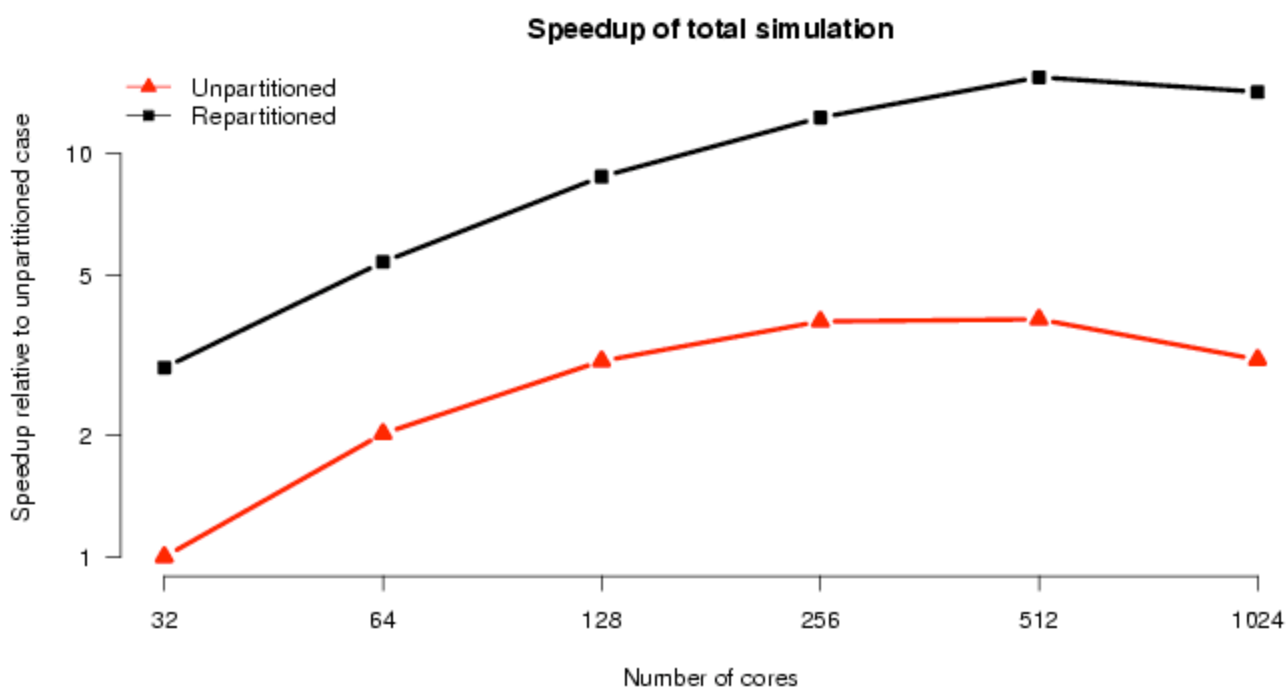

IO servers

- cf T. Edwards (this meeting)
- Dedicate small number of cores to output (hide latency)
- Scatter data from compute cores onto output cores
- Do mapping to canonical mesh on these cores

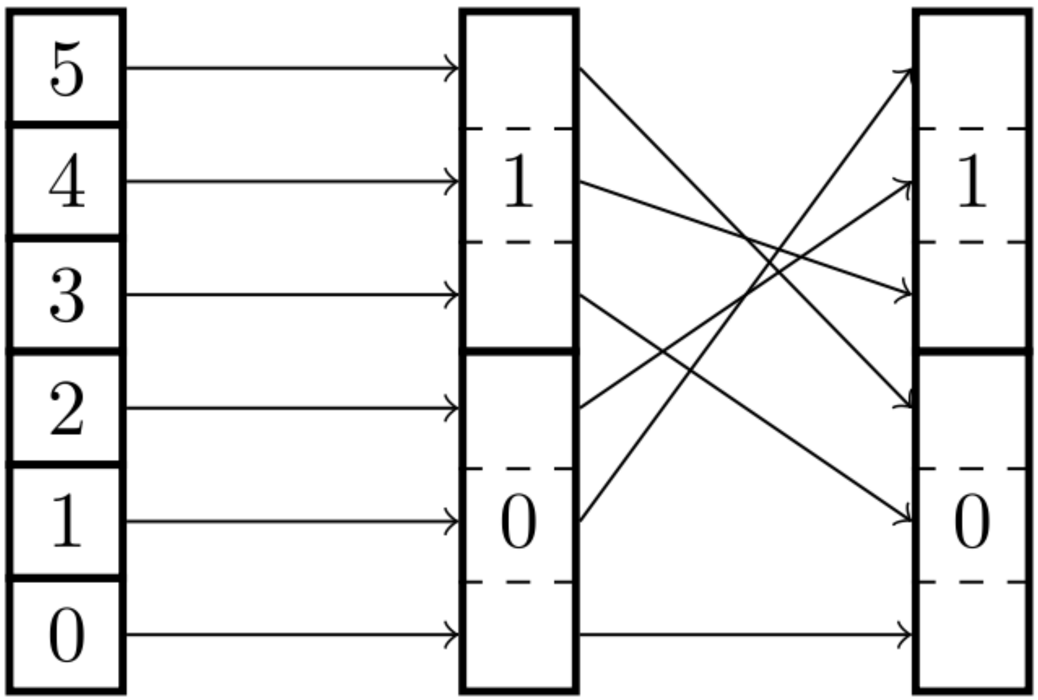

Simpler algorithms

- Can do explicit integration of parabolic PDE, rather than Krylov solve
- Just have to do matrix-vector multiply, no preconditioning or dot products to test convergence
- Timestep is smaller (stability), but each step is much simpler
- Wins in scaling *and* performance

Modelling a human heart

- 26 million unknowns, 150 million elements
- No elliptic PDE

Figure courtesy S.A. Niederer and G. Plank

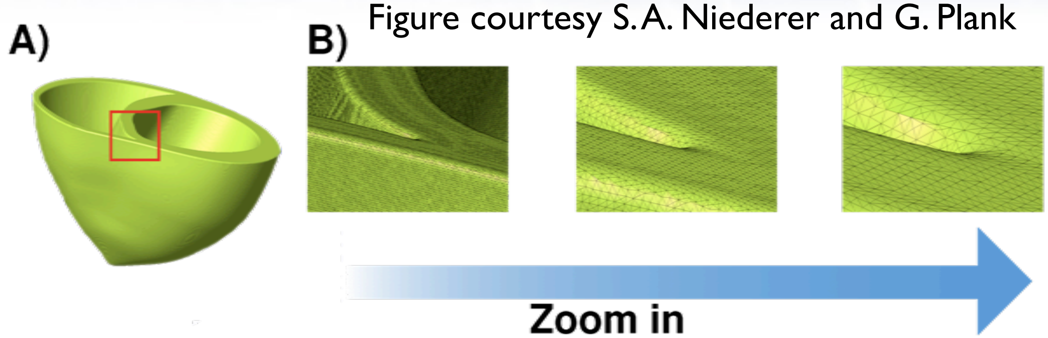

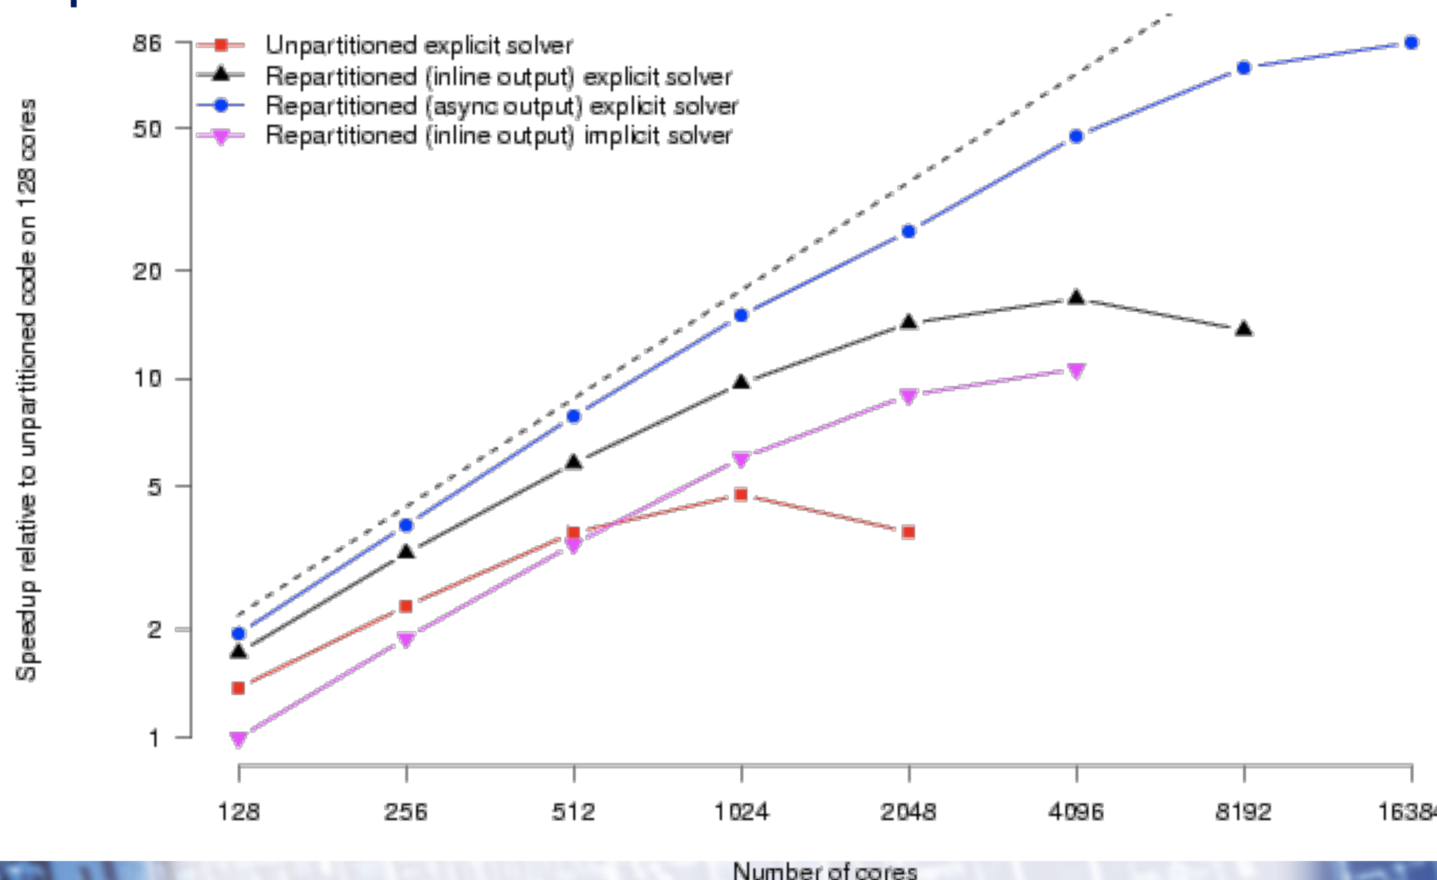

New science

- Simulations lag real time by factor of 280 (16k cores), previously 4300 (1k cores)
- 1 second activity takes 5 minutes, not 74. Speed-wise almost ready for deployment in pre-op planning scenarios.

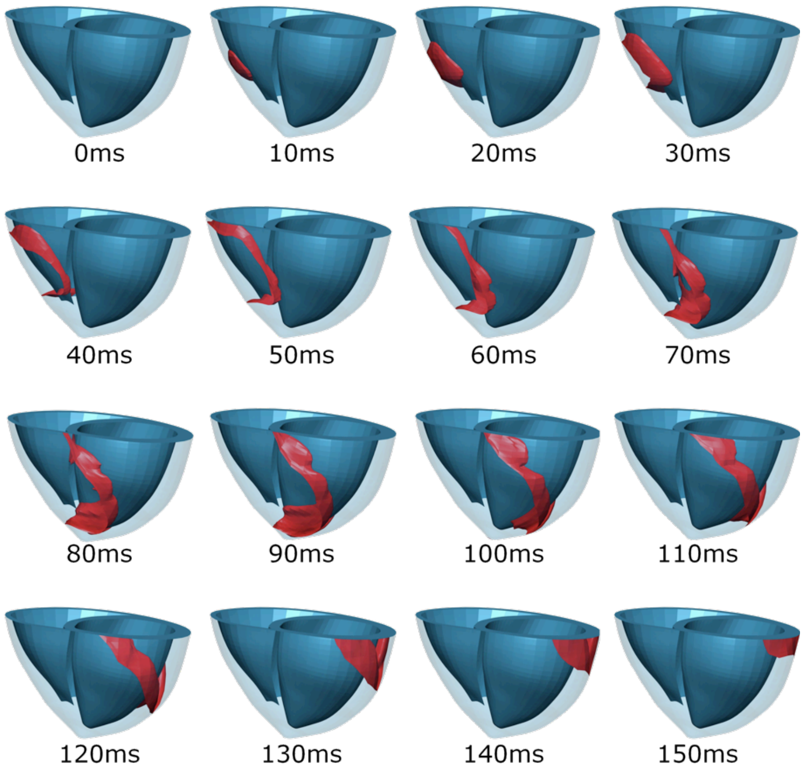

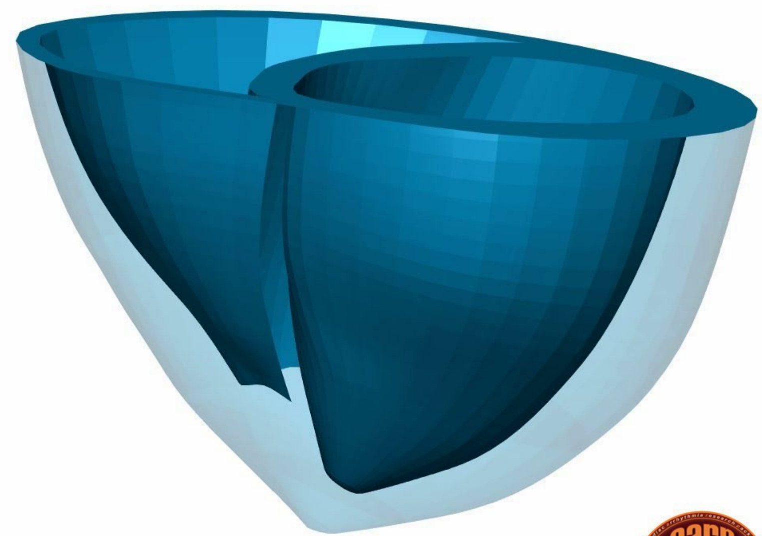

Niederer SA & Plank G

Conclusions

- Profile, and address each hotspot in turn
- Good partitioning is essential, but harder than regular grids
- Need to hide output latency. Can't do 16k parallel writes
- Revisit your algorithm choices

Acknowledgements

Work funded by a distributed CSE (NAG/HECToR/EPSRC) and DEISA WP9

Human heart data courtesy of S.A. Niederer and N. Smith (University of Oxford)

Rabbit heart data courtesy of M. J. Bishop and G. Plank (University of Oxford)

Questions?