Welcome to the EMiT Conference 2015

Benedict D. Rogers

University of Manchester

POLYHEDRON SOFTWARE
PROGRAMS FOR PROGRAMMERS

NVIDIA

OCF

intel

EAO
Doctoral Training Programme

hp

Harwell Imaging Partnership
Science & Technology Facilities Council

Background to EMiT Conference

In the past 15 years the hardware for HPC has changed enormously:

GPUs, FPGAs, hybrid chips, multi-core and massively parallel machines.

Energy consumption and maintenance have become key issues for **Green Computing**.

We now have a great range of potential architectures … but

In many cases, we as developers in universities cannot keep up with the appearance of new hardware (*application specific algorithms …*).

e.g. we release open-source software used over the world for Smoothed Particle Hydrodynamics (SPH), but we have spent 7+ years recoding on GPUs to make the method attractive to industry.

Aim of EMiT Conference

This raises an important question:

How do we:

- make the maximum use of the new hardware,
- keep ahead (not up-to-date) of technology emerging, and
- develop our numerical algorithms to future proof developments?

"The aim of EMiT is to bring together leading key figures in the computing communities, the end users of new software & hardware in industry and the vendors, to drive forward the use of emerging technology and get the most out of the extraordinary hardware in development."

EMiT COMMITTEE

2 Days of Talks

Sessions:

Novel Hardware 1 & 2

Low-power & Energy Efficient Computing

Real-time GPU Computing

Compilers & Portability

Applications

We will be drawing together key themes of the day and posing some awkward and provocative questions!!!

Afternoon Schedule

Keynote Speakers:

Prof. Stephen Furber, University of Manchester:
"SpiNNaker and the Human Brain Project"

Prof. Laura Grigori, Inria:
"Communication avoiding algorithms"

Dr John Linford, ParaTools:
"Kppa: A high performance source code generator for chemical kinetics"

Dr Filippo Mantovani, Mont Blanc Project:
"High Performance Computing based on mobile embedded processors"

MANCHESTER 1824

The University of Manchester

Manchester History

First **some history**: (Manchester is a **very humbling** place to work)

Manchester emerged as the **world's first industrial city**.

Manchester businessmen and industrialists established the **Mechanics' Institute** (Owen's College, UMIST, modern University of Manchester) to ensure their workers could learn the basic principles of science.

University can count **25 Nobel Prize winners** amongst its current and former staff and students

MANCHESTER 1824

The University of Manchester

Manchester History

MODERN ATOMIC THEORY

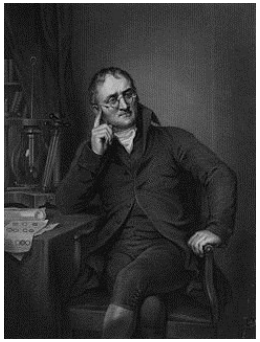

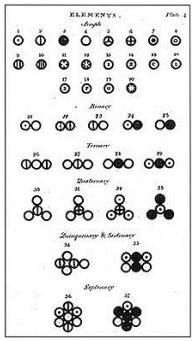

John Dalton

1st LAW of THERMODYNAMICS

James Joule

$$dU = \delta Q - \delta W$$

TURBULENT FLOW

Osborne Reynolds

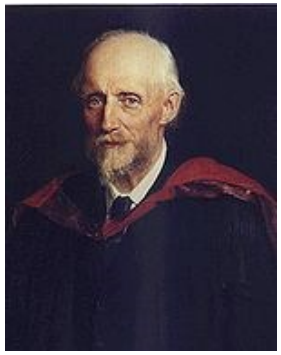

$$\mathrm{Re} = \frac{\rho V L}{\mu} = \frac{V L}{\nu}$$

JODRELL BANK

One of largest radio telescopes in world

WORLD'S FIRST TRAIN STATION

MOSI

WORLD'S FIRST PROFESSIONAL FOOTBALL LEAGUE

Mr Rooney

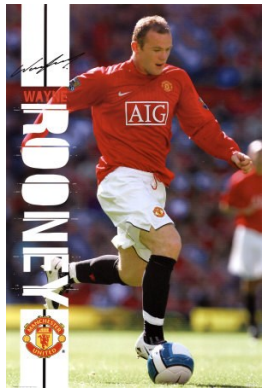

MANCHESTER 1824

The University of Manchester

Role of Manchester in Scientific Development

More Manchester Firsts :

COMPUTERS 1: The first computer with a stored programme and memory, nicknamed "Baby" or "Manchester Mark 1", was developed at Manchester University in 1948 by Profs Tom Kilburn and Fred Williams. Changing a program would take minutes rather than days.

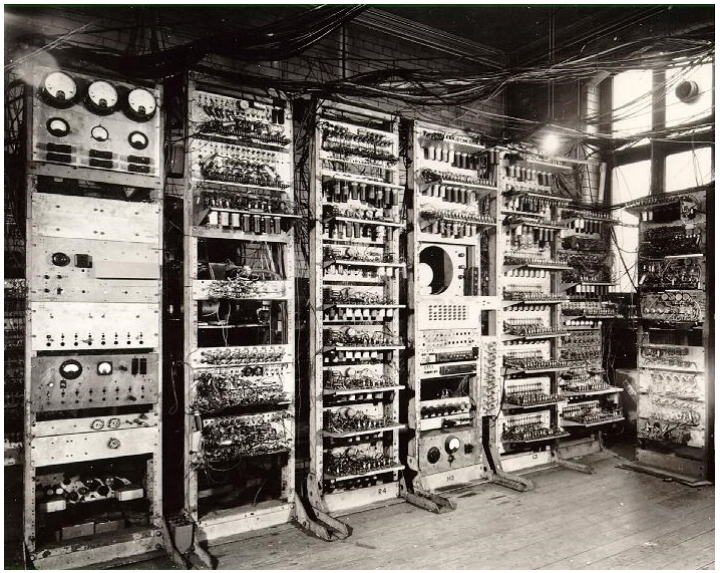

COMPUTERS 2: Alan Turing (1912 – 1954) played a significant role, also developed the Turing Test for artificial intelligence

MANCHESTER 1824

The University of Manchester

Latest Nobel Prize winners

2010:

GRAPHENE: Thinnest supermaterial in the world

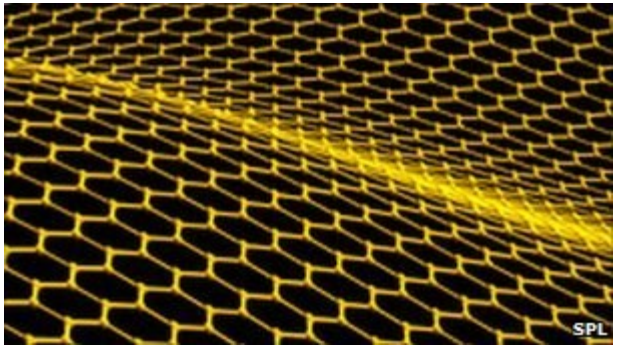

Profs. Novosolev & Geim

Computing Facilities at Manchester

Lots of small machines for teaching and exploratory research:

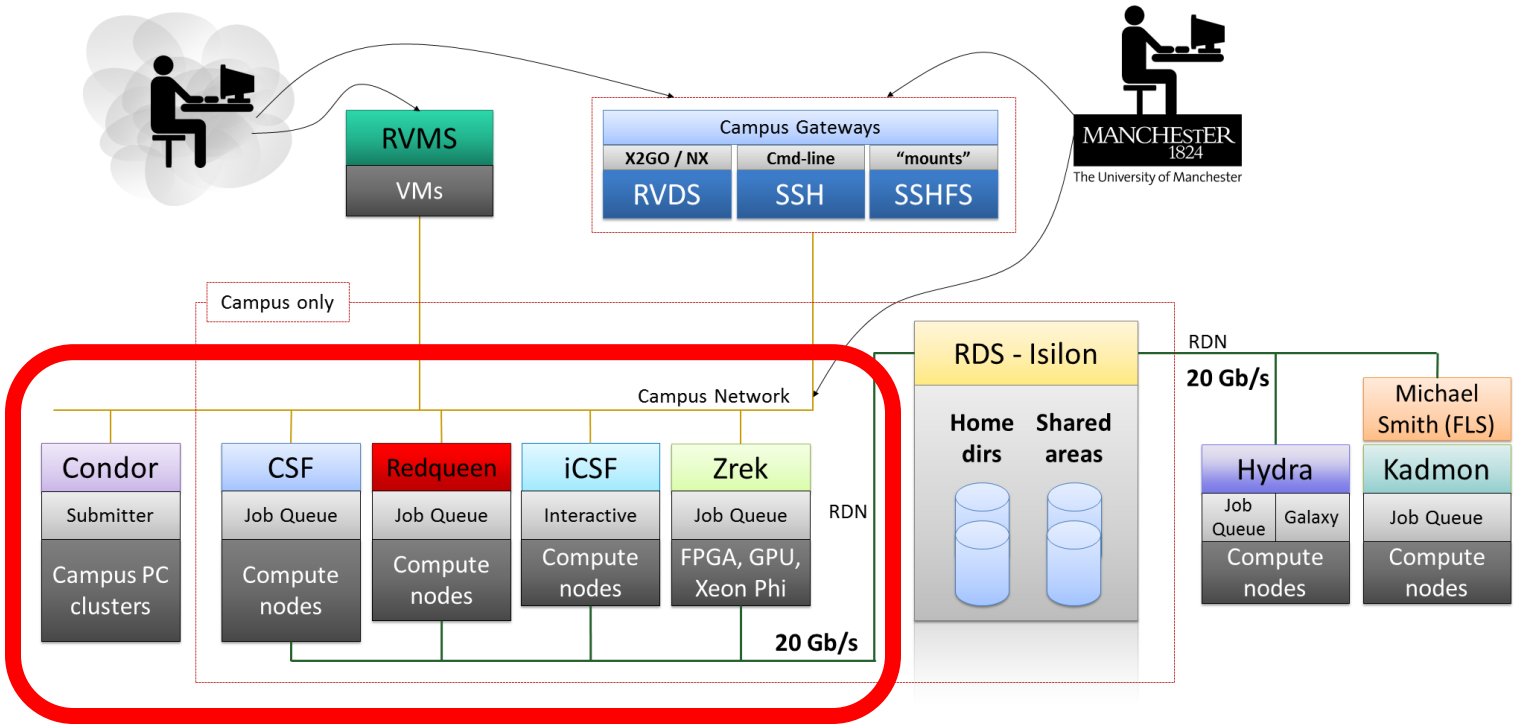

Key web sites:

- http://gpu.manchester.ac.uk
- http://ri.itservices.manchester.ac.uk
- http://wiki.rac.manchester.ac.uk/community/Courses

Emerging Technology Facilities at Manchester

"Computationally Intensive Research Ecosystem" (CIRE)

- Multi-GPU cluster (NVIDIA K40s)
- Xeon Phis
- FPGAs

CSF: Conventional 5000 core cluster

With high memory nodes (256 GB and 2 TB RAM)

For Visualisation: **iCSF Graphics/Viz Capability**

GPU Club (Michael Bane)

Image-Based Modelling Club

Finite Element Club

Multi-scale Club

Scientific Programming Club

Since 2013, Manchester is a CUDA Research Centre (CRC)

NVIDIA

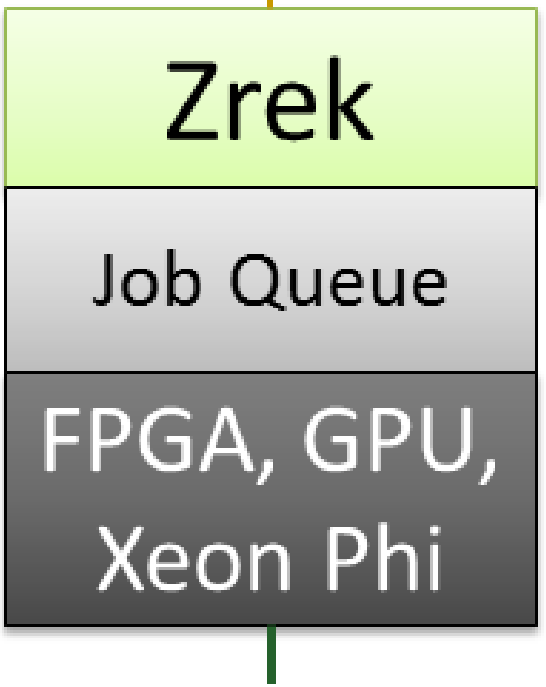

Computing Facilities outside of Manchester

Access to National and Regional Facilities:

N8-HPC

Hartree with STFC (Daresbury)

ARCHER – national HPC service

EMERALD – large multi-GPU cluster

Work at Manchester: Applications

MANCHESTER 1824

MANCHESTER 1824

The University of Manchester

GPUs Enabling Biological Research

Andrew Almond & Ben Sattelle

Manchester Institute of Biotechnology

Almond & Sattelle perform simulations of complex sugars (carbohydrate polymers), which mediate inflammation and cancer via their interactions with protein signalling molecules (Figure 1) and comprise the basis for energy storage in plants. The research aims to improve biological understanding, it focuses on accurate modeling of molecular 3D-shape and has applications ranging from development of new medicines to food sustainability.

Figure 1: Structure of Human Extracellular Heparan Sulfate Modelled Using GPUs

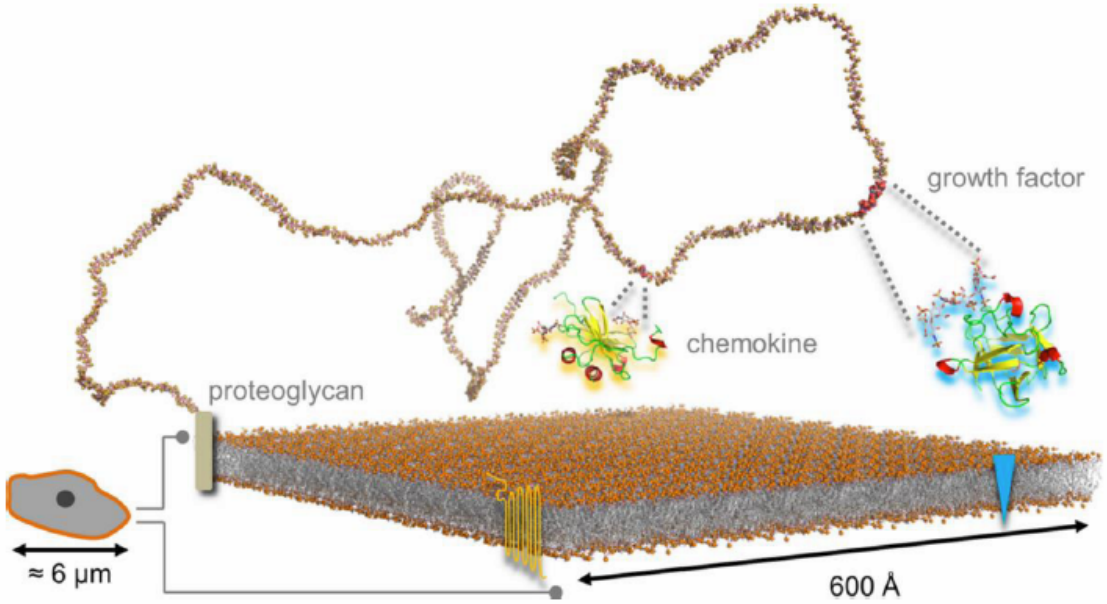

MANCHESTER 1824

The University of Manchester

Automating HPC Runs of WRF and WRF-Chem for the ManUniCast Project

ManUniCast is a project run by the Centre for Atmospheric Science (CAS) at The University of Manchester. The aim of this project was to build a real-time weather and air-quality forecasting website, which could be used as an educational tool by students and the general public. This involved obtaining initial conditions from various sources, running a computer simulation and then preparing appropriate output.

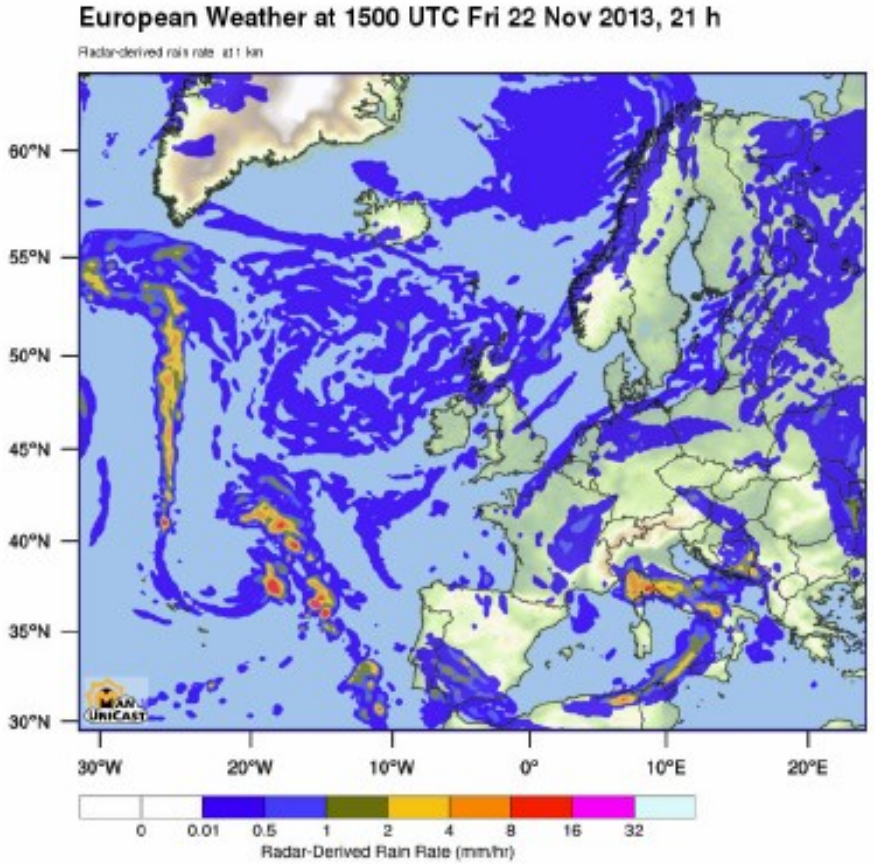

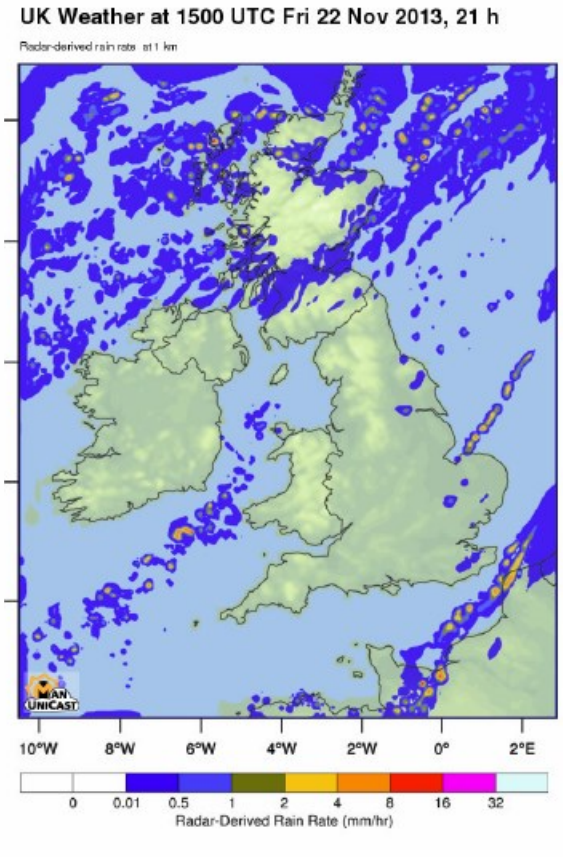

MANCHESTER 1824

The University of Manchester

Advanced Processor Technologies (APT)
Prof. Furber

SpiNNaker is a novel computer architecture inspired by the working of the human brain.

A SpiNNaker machine is a massively parallel computing platform:

- Neuroscience
- Robotics
- Computer Science

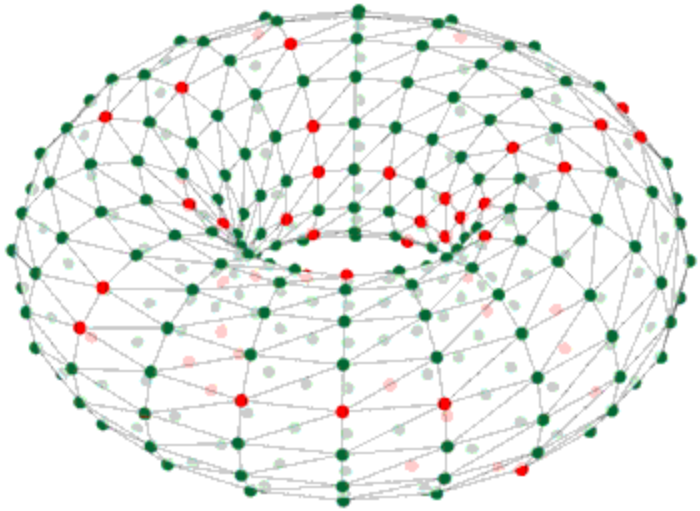

Smoothed Particle Hydrodynamics (SPH) free-surface Applications

MANCHESTER 1824

Original Motivation for SPH

- **Free-surface flows** are rarely singly connected, e.g. beaches & wave energy devices

Breaking waves on beaches

(Photo courtesy of F. Raichlen)

Very complex Multi-phase Multiscale problems

Wave Energy Devices: Manchester Bobber

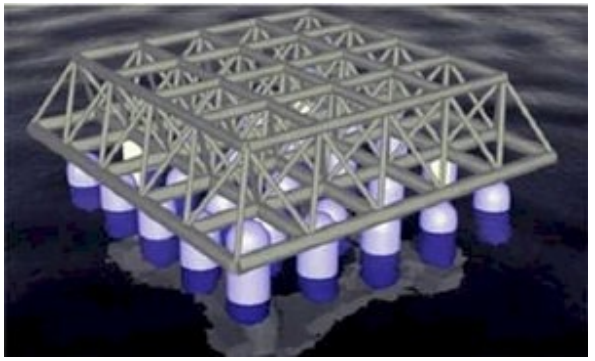

Smoothed Particle Hydrodynamics (SPH)

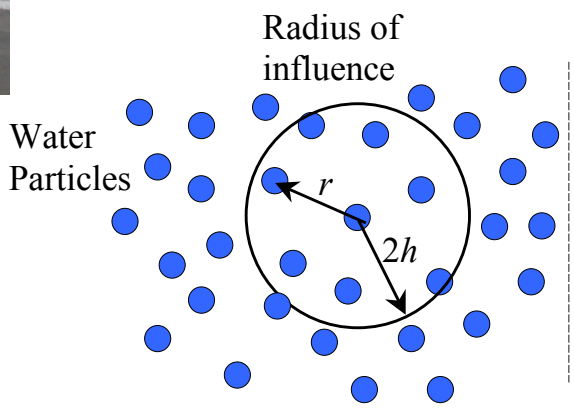

SPH: Is one of the few simulation techniques that can handle these situations

Classical SPH Formulation Example

2011 Japanese Tsunami

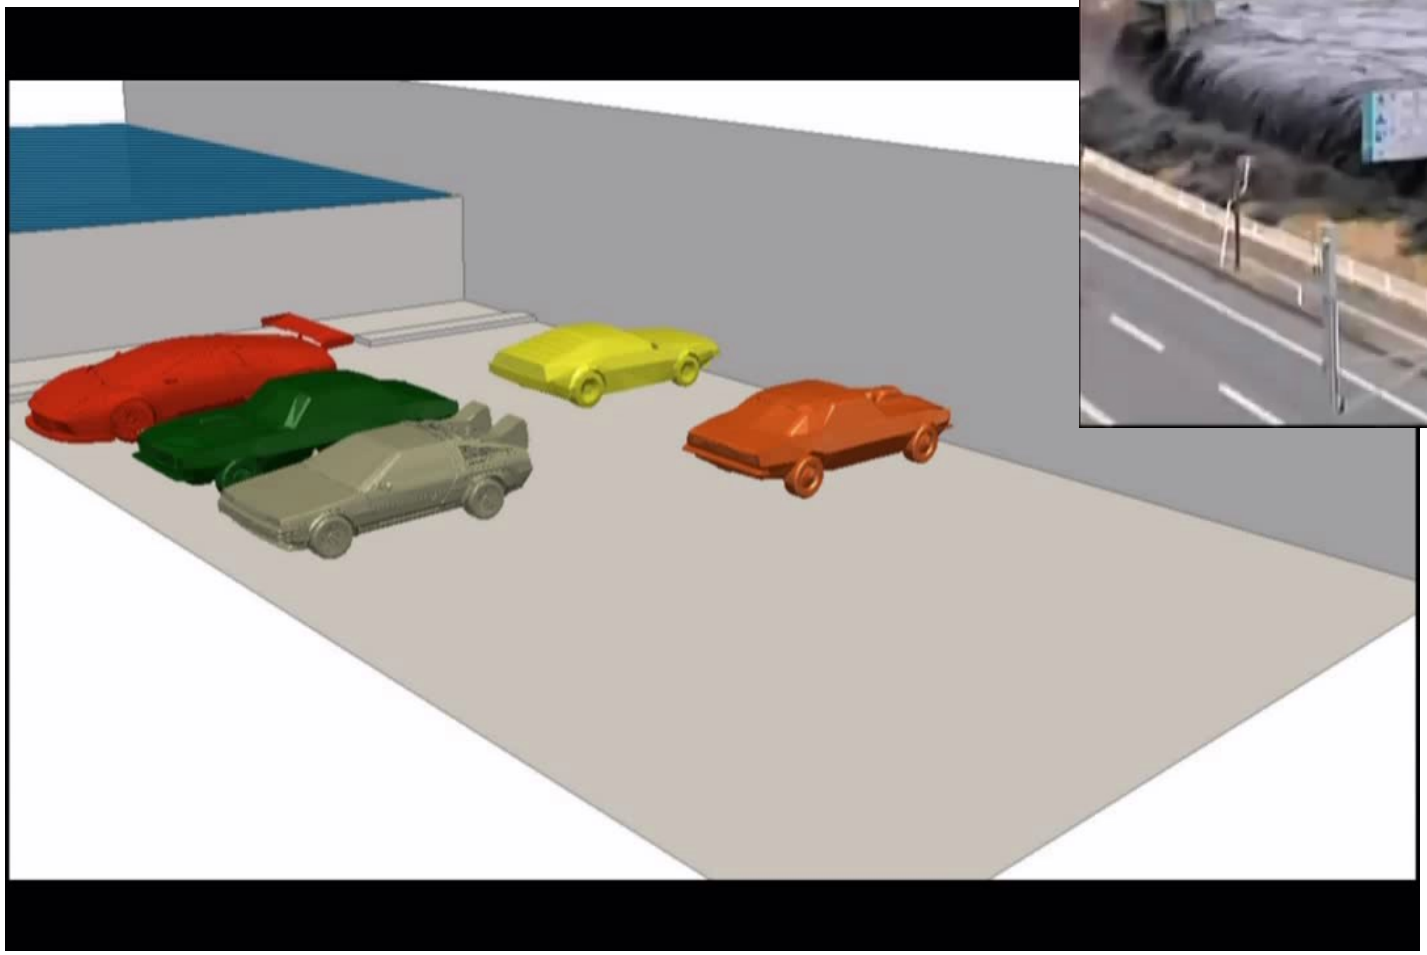

cpu gpu
DualSPHysics

Single GPU code which can handle 10millions particles (Crespo *et al*. 2012)

Multi-GPU SPHysics

- Dual-SPHysics (cpu & gpu) – SPEEDUPS of **400+** for a Desktop machine = Supercomputer of 1000's cores
- **Valdez-Balderas et al. (2012)**

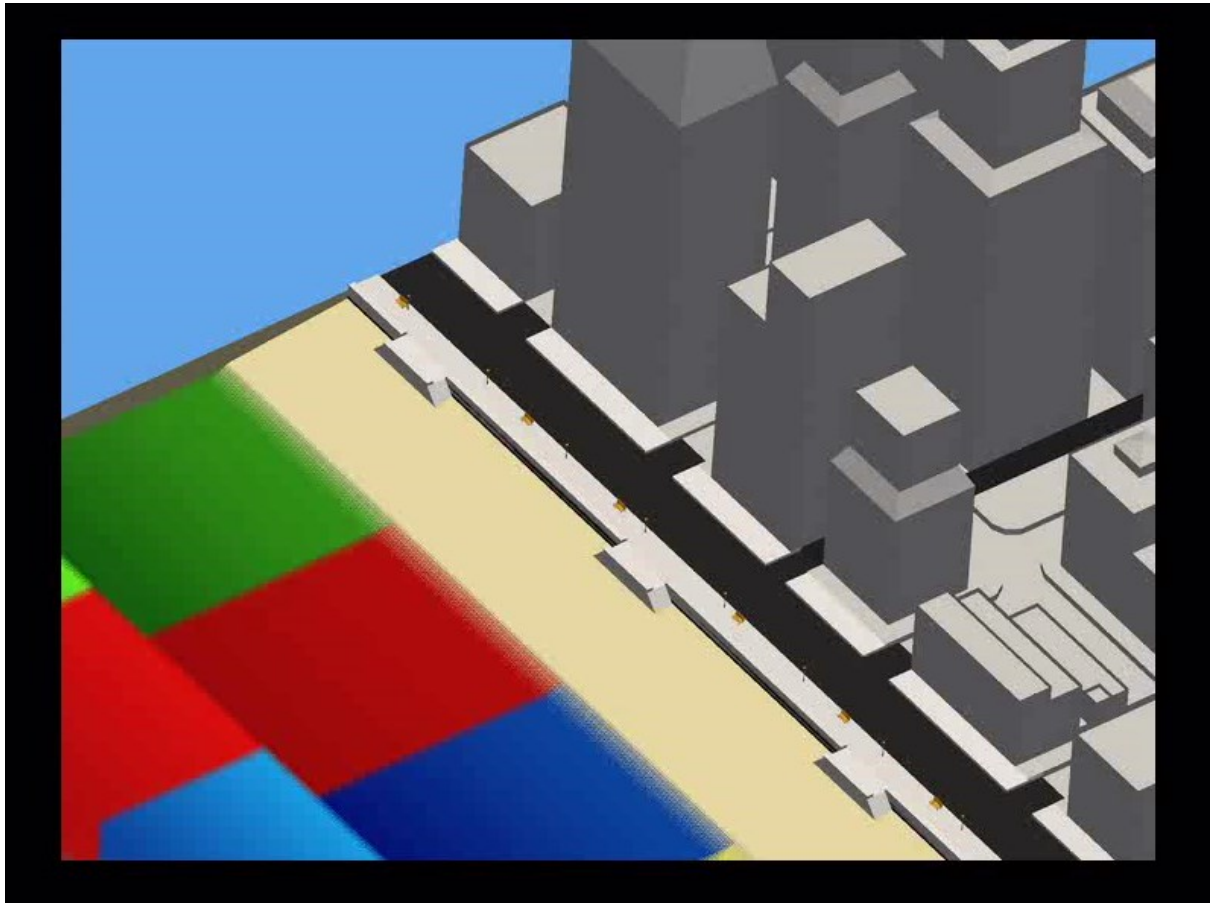

cpu gpu
DualSPHysics

4 GPUs	2.1 hours
1 GPUs	5.7 hours
1 CPU (4 core)	3.5 days
1 CPU (1 core)	2 weeks

5.3 million particles

For 3-D problems (convergence & application), this requires 100s millions of particles!!!

SPH 1 billion Particles!! MultiGPU DualSPHysics

Large wave interaction with oil rig using 1 BILLION (10^9) particles

[Click here for Movie 1]

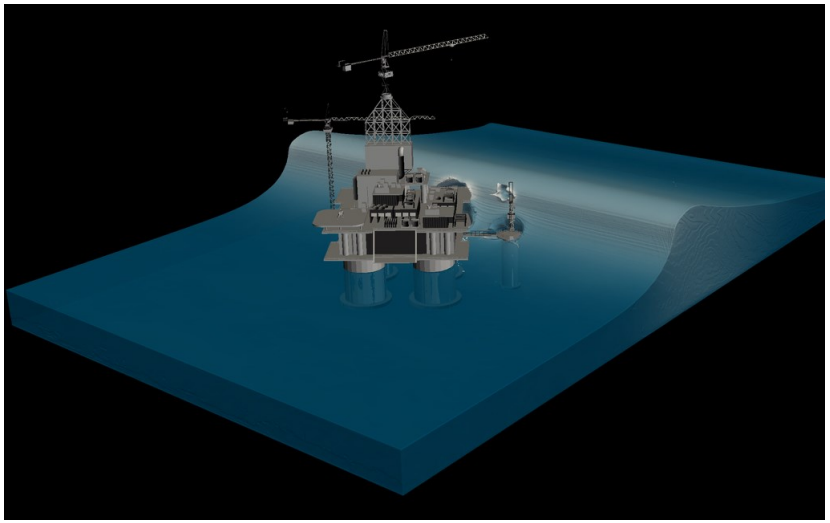

Dominguez et al. (2013) Computer Physics Comm.

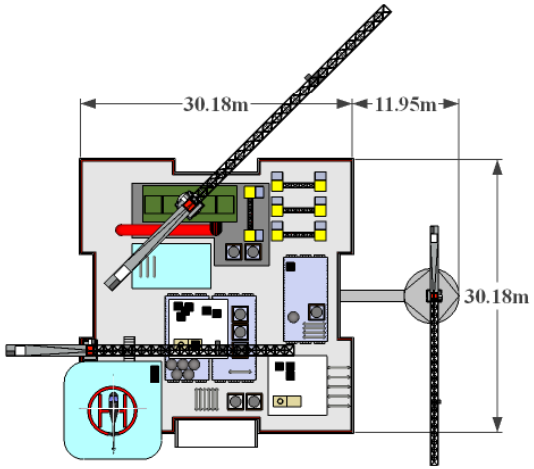

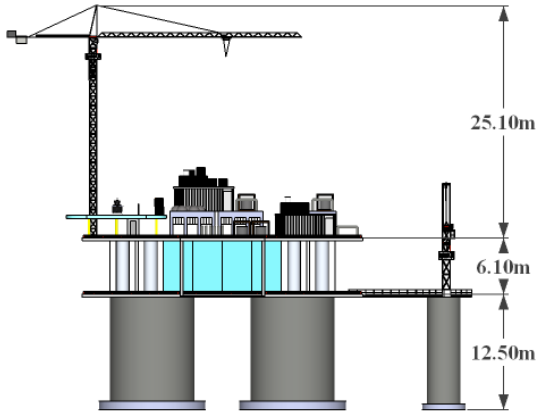

We have post-processing tools efficient enough to visualise the simulation of more than 1000M

dp = 6 cm
np = 1,015,896,172 particles
nf = 1,004,375,142 fluid particles
physical time = 12 sec
# of steps = 237,342 steps
runtime = 91.9 hours

using 64 GPUs Tesla 2090 of the BSC

SPH 1 billion Particles!! But …

This required **6+ years of recoding** by a team of researchers based at 3 universities.

We didn't know (and still don't) how long GPUs will be a viable choice for architecture selection.

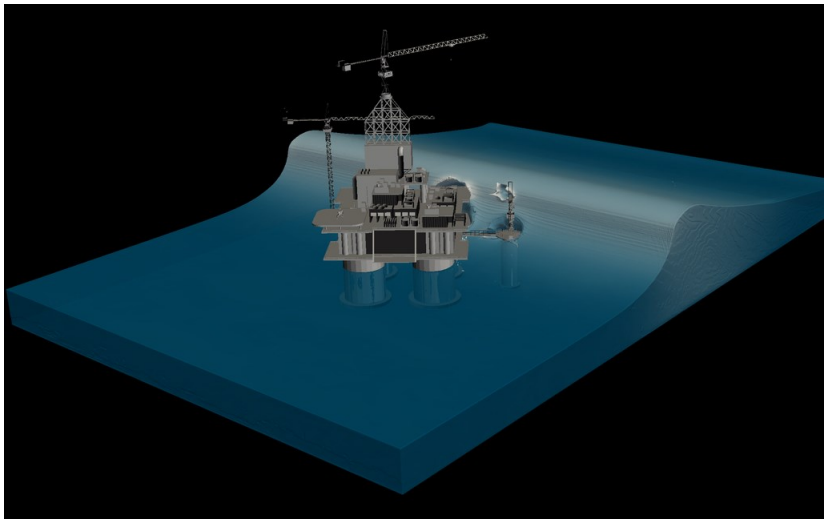

Is there a smarter way to do this?

Better collaboration?

Better use of facilities?

Encourage use within Industry?

We don't want to recode again …

… EMiT 2015 Conference

This afternoon's Panel Discussion!

Morning Schedule

Time	Speaker	Organisation	Topic
09:45-10:30	Prof. Stephen Furber	The University of Manchester	SpiNNaker and the Human Brain Project
Break (10:30-11:10; 40mins)			
Session 1: Novel Hardware 1 (11:10-12:10; 1hr)			
11:10-11:30	Wim Vanderbauwhede	University of Glasgow	Exploring Deep, Hierarchical Pipelines and High-level Synthesis on FPGAs for Accelerating a Legacy Convection Model
11:30-11:50	Peter Düben	University of Oxford	Reduced numerical precision in weather and climate models
11:50-12:10	Prabu Thiagaraj	The University of Manchester	FPGA processing for High Performance Computing
Lunch (12:10-13:10; 1hr)