Keele University

NAM 2014

IPMU INSTITUTE FOR THE PHYSICS AND MATHEMATICS OF THE UNIVERSE

Knowledge Exchange Session

HPC-collaboration with industry on future computing platforms

Raphael
HIRSCHI

http://www.astro.keele.ac.uk/shyne/

Plan

- Introduction
- Scientific goals
- ERC project & simulation frameworks
- Computing platforms
- Numascale architecture
- Benefits of collaboration between scientists and HPC industry

Astrophysics

Goal: understand the properties and evolution of the universe and its constituents

BIG BANG

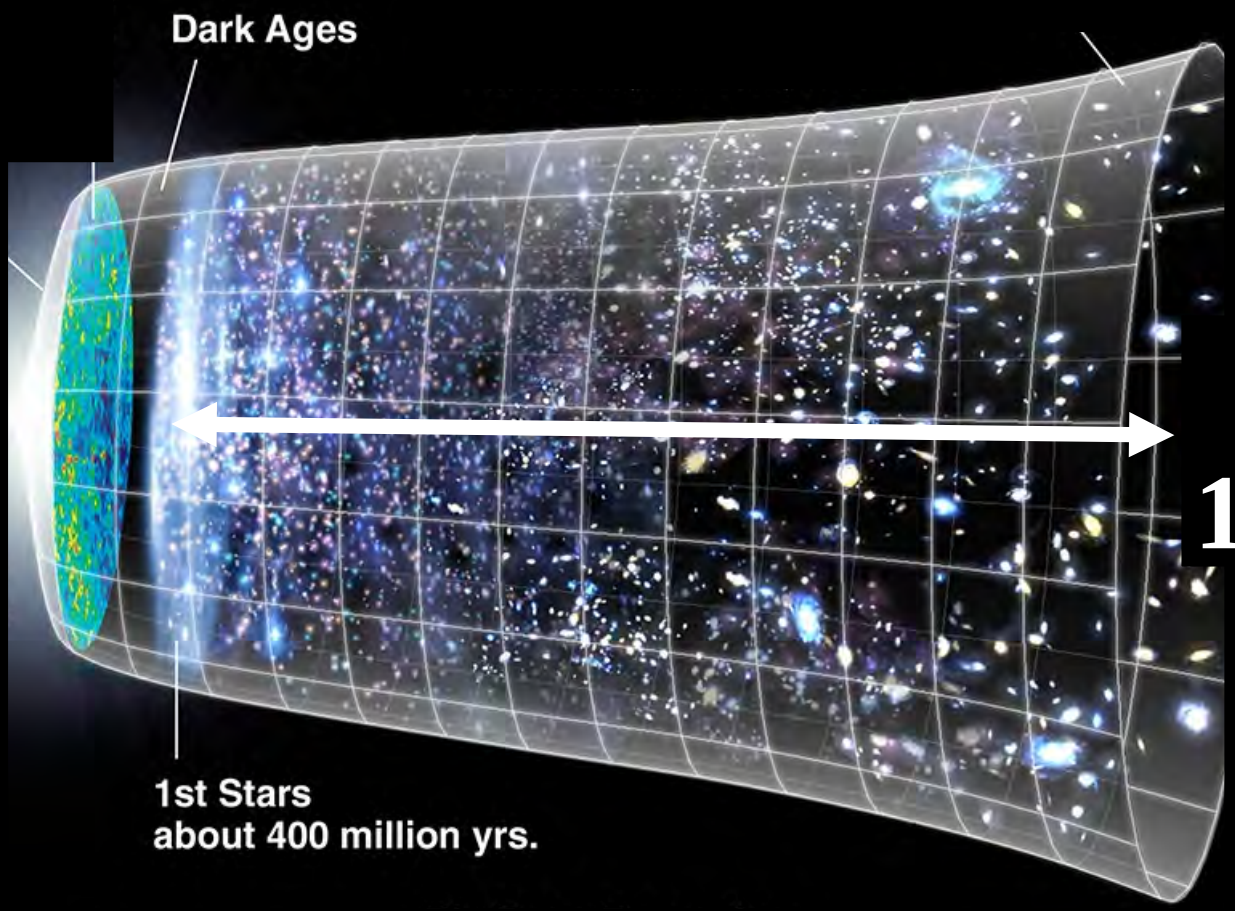

NOW
13.7 Gyr

Keele University

PI: Raphael HIRSCHI

IPMU INSTITUTE FOR THE PHYSICS AND MATHEMATICS OF THE UNIVERSE

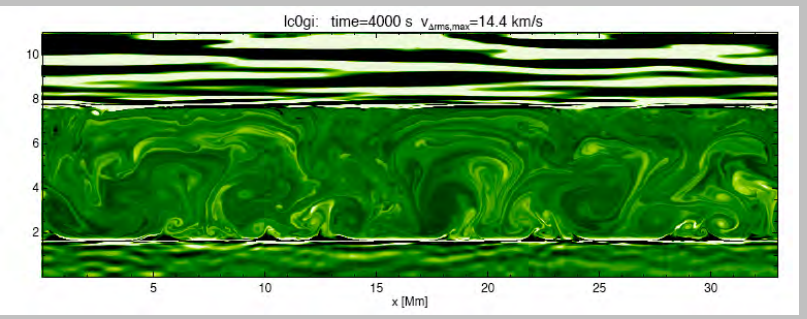

©Herwig

erc
European Research Council
Established by the European Commission
Supporting top researchers from anywhere in the world

SHYNE:

Stellar HYdrodynamics Nucleosynthesis & Evolution

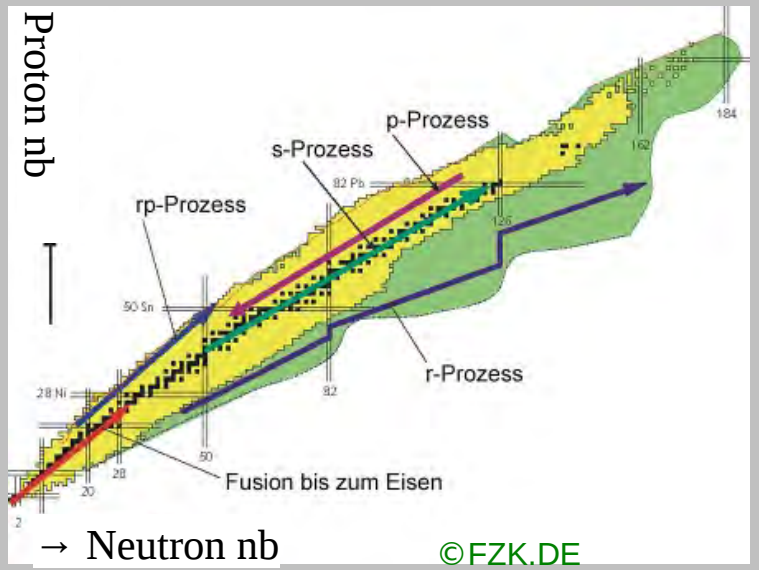

©FZK.DE

- Key facts:
- ERC starting grant
- 5 year
- EURO 1.4M
- 2PDs+2PhDs
- 1000+ CPUs cluster

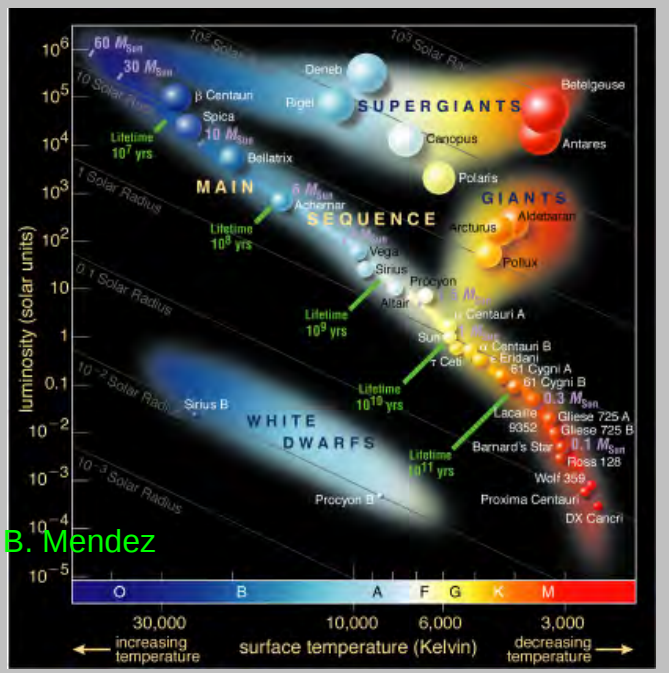

©B. Mendez

Stellar HYdrodynamics Nucleosynthesis & Evolution (SHYNE) Project

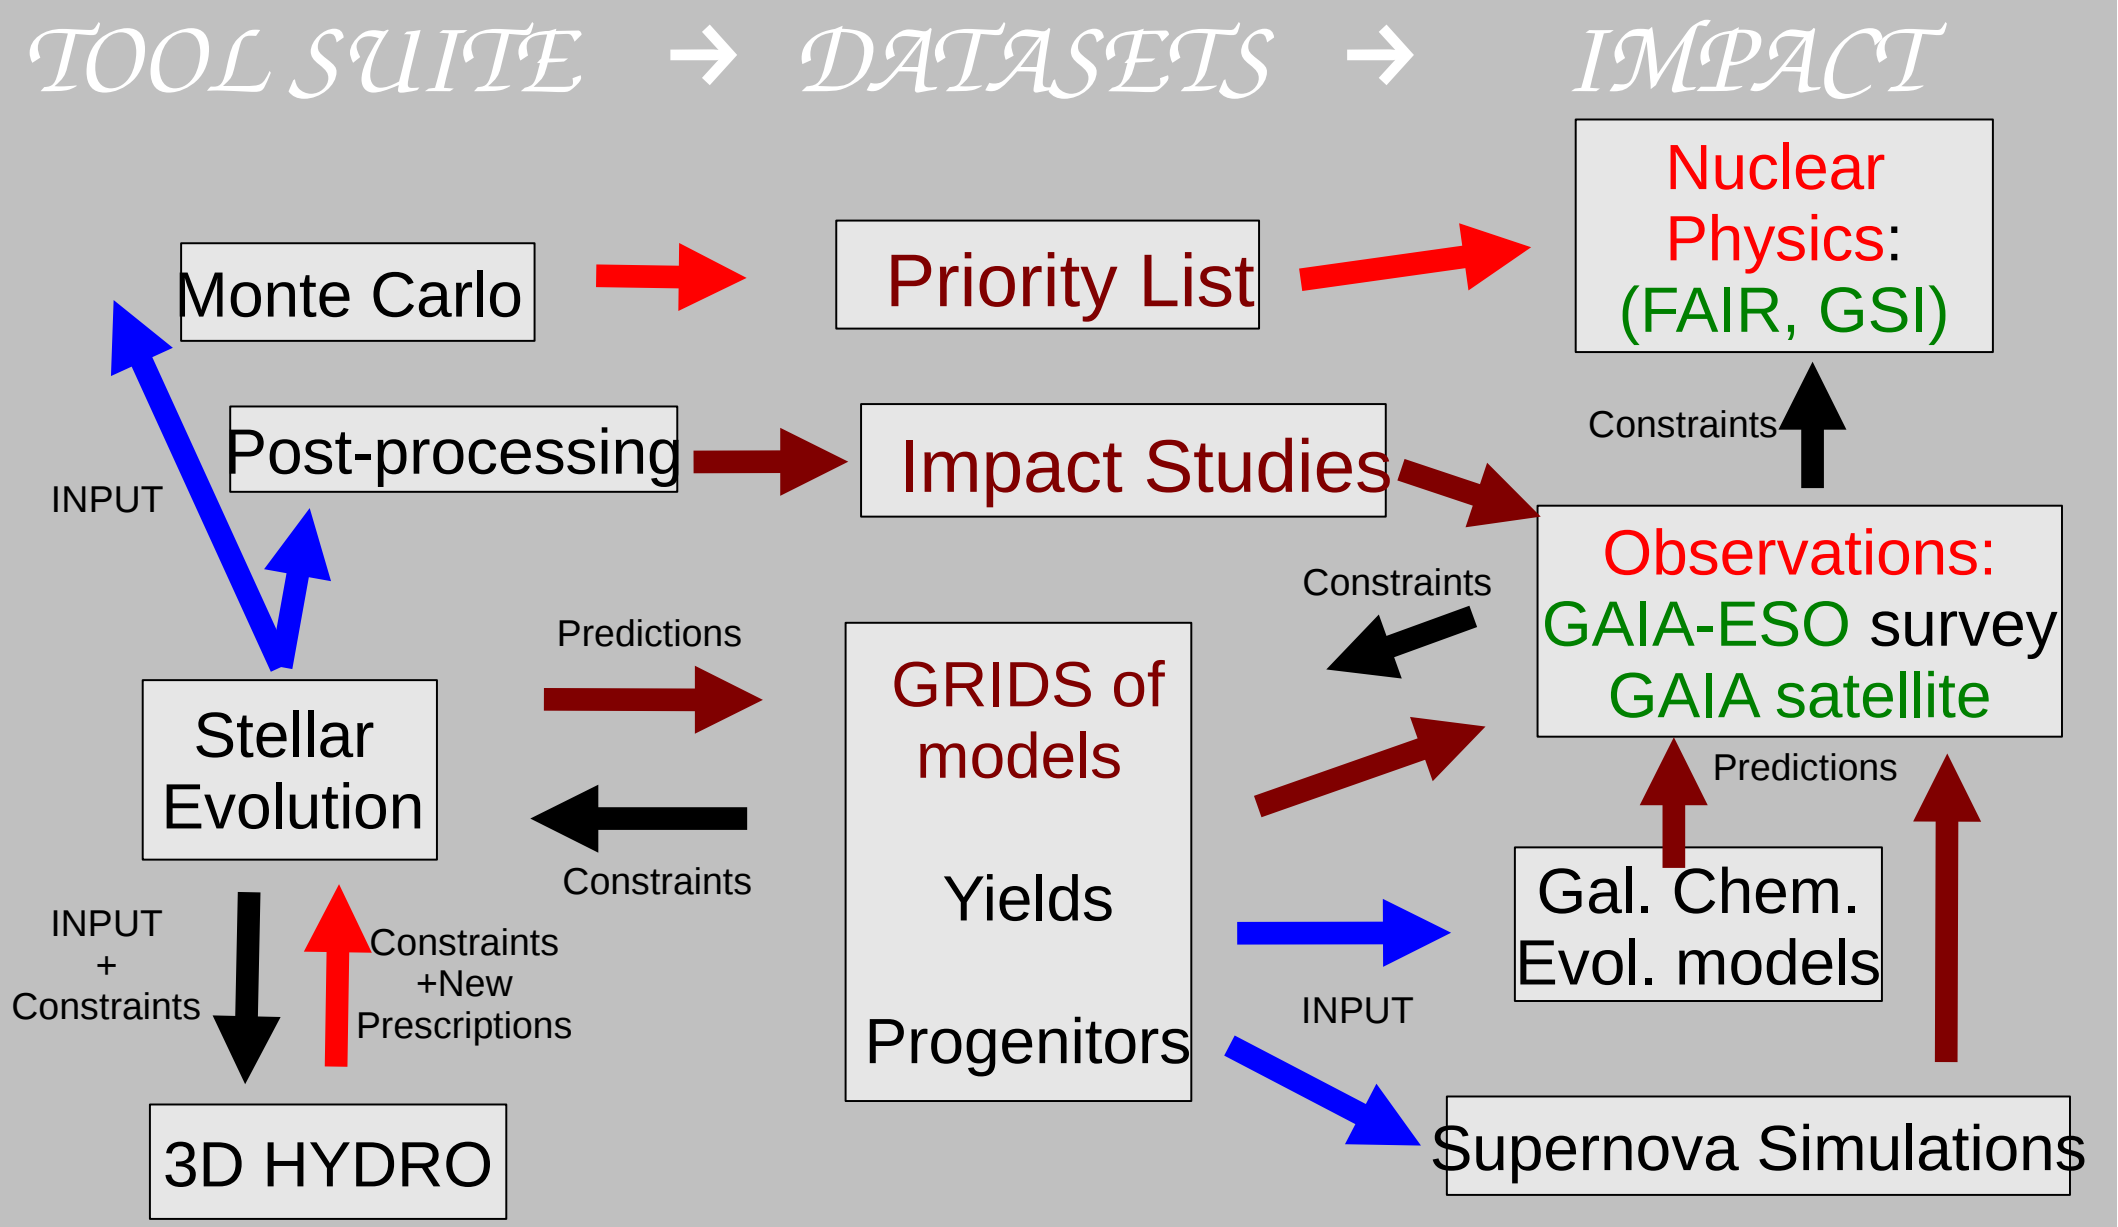

- Efficient pipeline: nuclear/hydro/astro

Hydrodynamic Simulations Framework

3D hydro simulations for key phases/processes

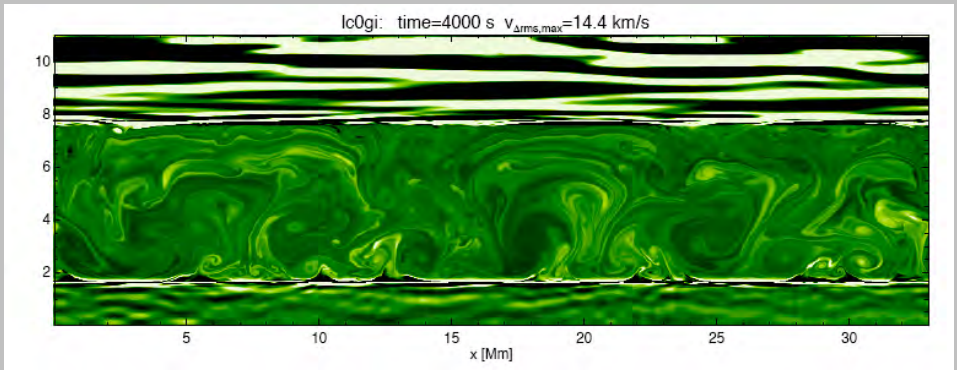

Herwig et al 06

Long-term evolution in 1D

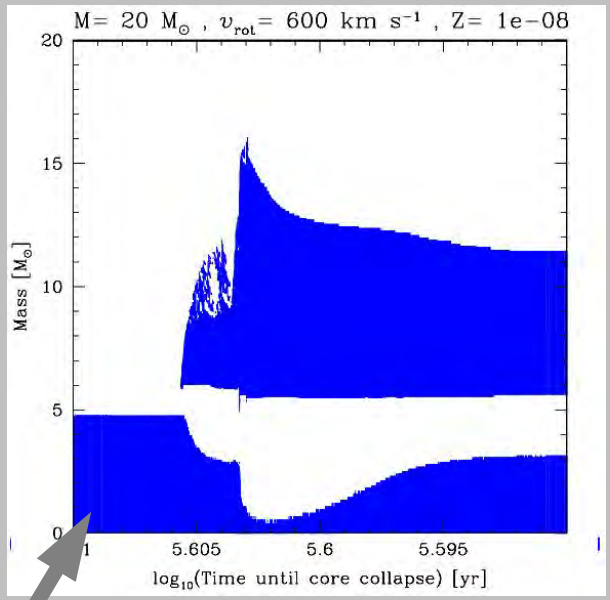

e.g. Hirschi 07

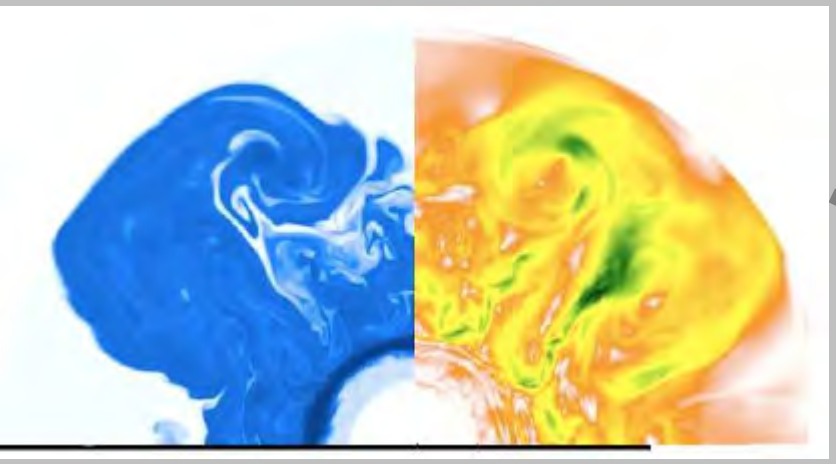

e.g. Arnett & Meakin 2011
Mocak et al 2011, ...

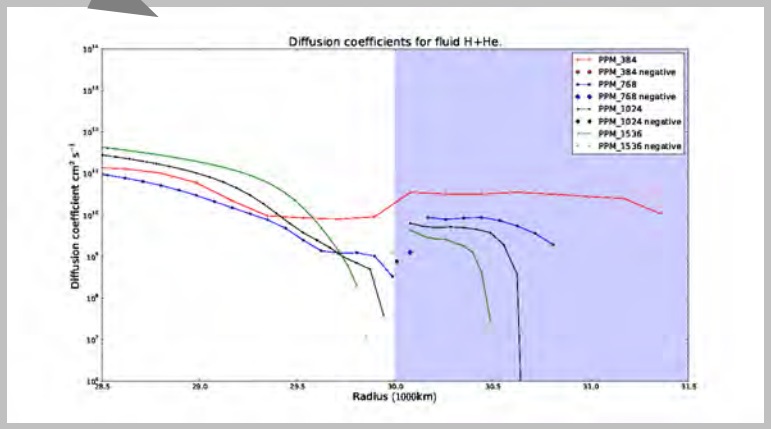

Meakin et al 2009 ; Bennett et al in prep

Determine effective diffusion (advection?) coefficient

Nucleosynthesis Simulations Framework

stellar evolution with minimal network for energy production

↓

multi-zone nucleosynthesis post-processing and mixing with adaptive, complete network for all nucleosynthesis

↓

all abundance data at all times in all locations inside the star

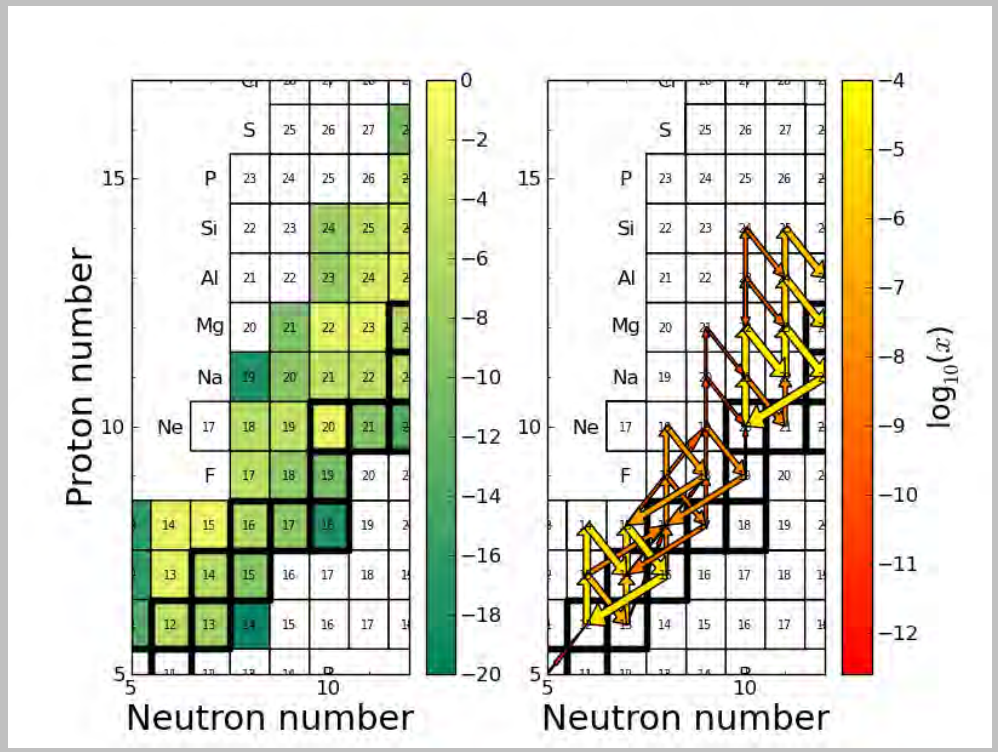

Example of Nova nucleosynthesis simulations (Pignatari & Herwig 2013, NPN), see www.nugridstars.org.

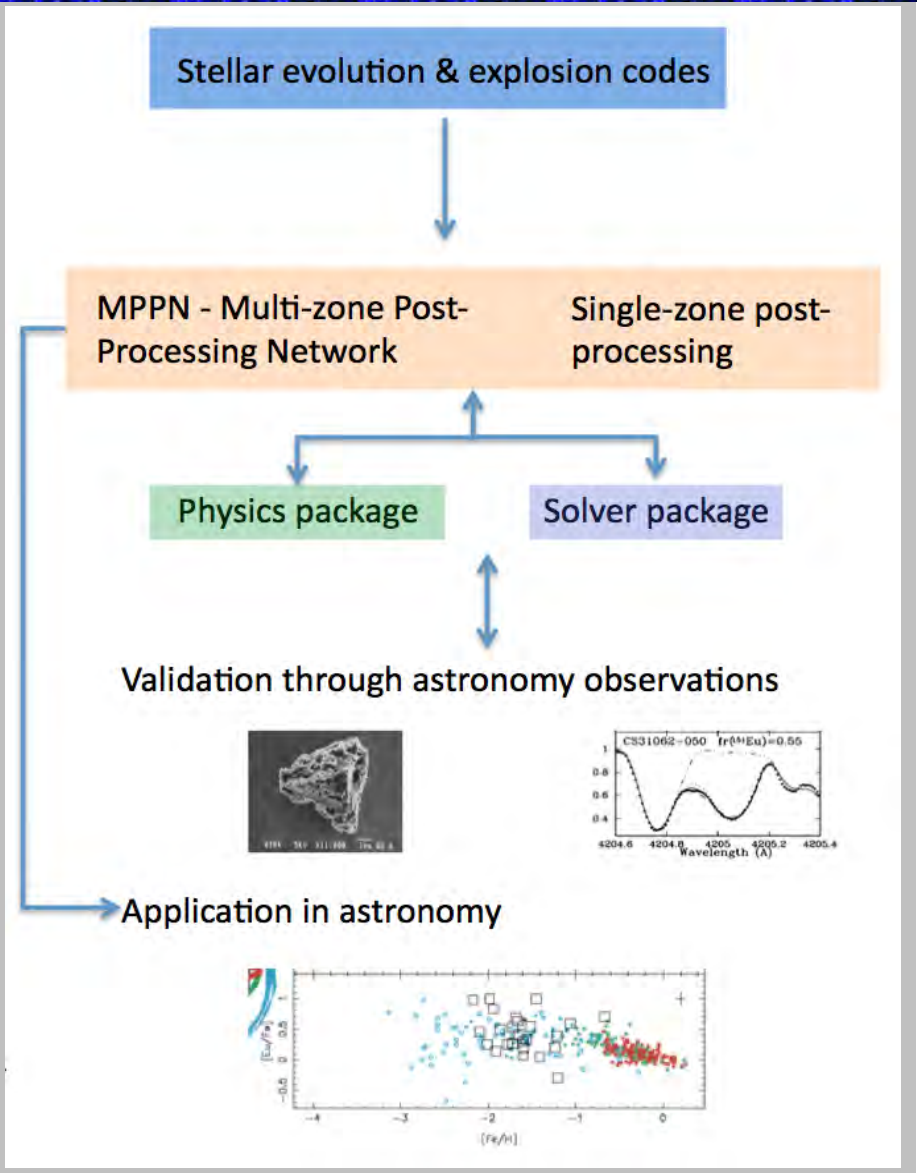

Guidance for future observation goals. E.g., **ASTRONET**

General Challenges and HPC Role

- SHYNE challenges are representative of challenges faced in a wide range of academic fields and Industries.

- How to harness the ever increasing computing power to tackle these challenges?

- In which direction is the HPC field going?

Computing Platforms: CPUs

- CPU speed has reached a plateau

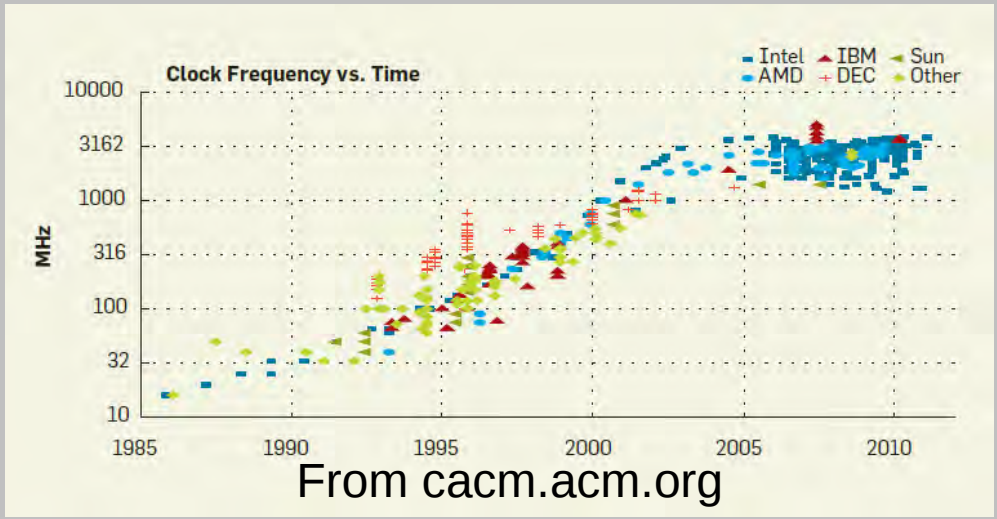

From cacm.acm.org

AMD:

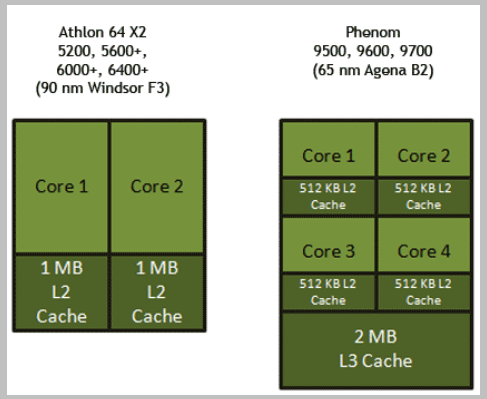

Intel:

http://www.tomshardware.com/reviews/phenom-athlon-core-scaling-compared,1759-2.html

- So increase in computing power only using multiple cores
- Good scalability BUT
- Need to modify existing codes (MPI)
- Limitation due to communication
- Power consumption

Computing Platforms: GPUs

- GPUs take distributed memory architecture to higher level
- Number of GPU- cores about 100x larger than regular CPU cores
- Power consumption much lower than CPUs

BUT

- RAM very limited
- Programming is not straightforward (although this is improving): need to re-design codes entirely?

Computing Platforms: Large Shared Memory Systems

Should we still try to build larger shared memory systems?

- Many scientific applications require large RAM
- Expertise and manpower needed to transport codes to multi-core/GPUs platforms
- Largest clusters may become inefficient without shared memory islands

SHYNE Computer Cluster

- The cluster comprises a total of 1056 AMD-based CPUs with 2GB RAM per CPU.

The main specifications are the following:

- 288 cores Numascale
  - 8 x 3 socket servers (each socket has 12 CPU cores)
  - Single memory image 576GB
  - Single operating system image
  - Numascale inter-connect
- 768 cores QDR Infiniband
  - 12 x 4 socket servers (each socket has 16 CPU cores)
  - 1 O/S per server
  - Distributed memory image 128 GB per server
- Unified cluster management for both architectures
  - IBM Platform HPC
  - 2 LSF queues Numa and IB
  - Numa nodes visible as single machine with 288 cores and 576 GB RAM
- Dedicated water cooled environment up to 30kW in 1 rack

Numascale Shared Memory Solution

numascale Norwegian HPC company

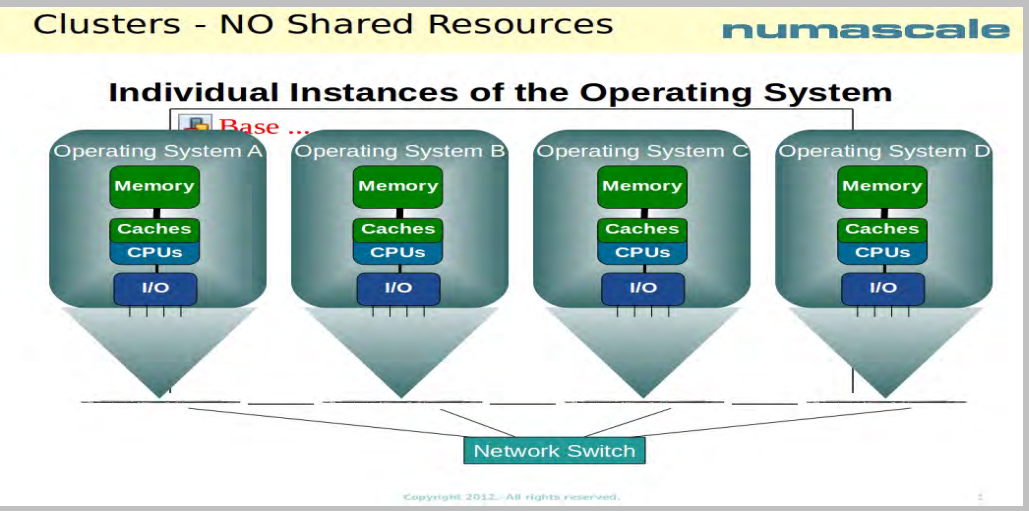

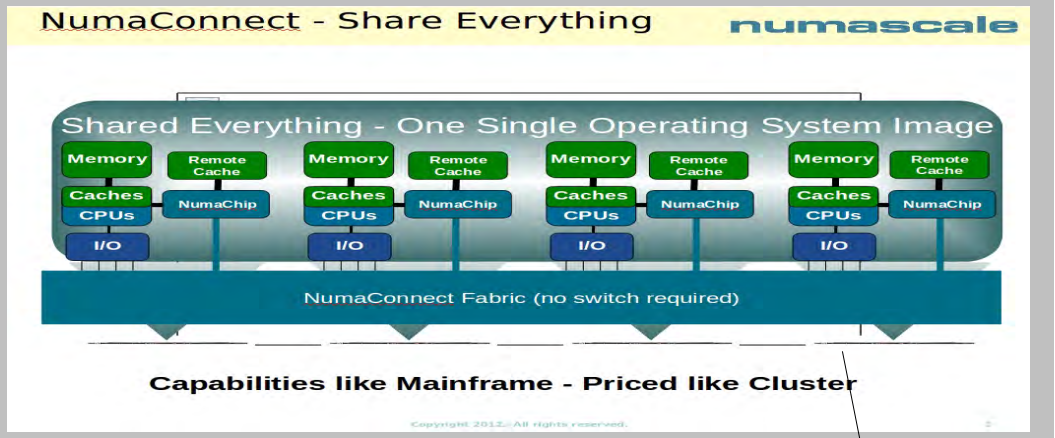

FROM: distributed memory clusters

TO: scalable shared memory clusters

For the same cost!

- Super-desktop: single OS, 288 CPU cores, RAM 576 GB
- No need to adapt existing codes
- Results available soon … Watch this space!

2D-torus

Benefits of Collaborating with Industry

- Improved optimisation of codes for scientists
- Training of next generation
- Knowledge transfer between HPC industry and user community
- Performance feedback for manufacturer
- Shaping the design of future largest computer systems together

  e.g. “Towards Petaflops Capability of the VERTEX Supernova Code”, Marek et al, arXiv1404.1719