

Astro, HEP, and plasma Data and computing needs

Stefan Krieg-Venghaus | Jülich Supercomputing Centre







Present usage: US

	0	55	110	165	220	275
Material Science	266					
Fusion Energy	212					
Chemistry	160					
Lattice QCD	137					
Climate	134					
Geosciences	61					
Astrophysics	53					
Life Sciences	46					
Combustion	31					
Accelerator Physics	29					
Mathematics	20					
Nuclear Physics	20					
Computer Science	18					
High Energy Physics	14					
Environmental Science	4					
Engineering	2					
Nuclear Energy	1,3					
Other	<1					

NERSC 2013

DoE INCITE 2014+2015 allocations





Present usage: PRACE

PRACE 2014 (calls 8&9) allocations



04.02.2015











Astro: science case

- Dark energy, dark matter
- Evolution of the universe and its structures:
 - Galaxies
 - Stars
 - Planets
- Study AGNs, black holes etc.
- "Map the sky"
- Etc.

Models(theory) and observations rely heavily on available computing and software infrastructure



Astro: structure of the universe

- Simulations:
 - Substantial memory requirements, O(PB) to match experiment
 - Realistic simulations of galaxies need Exascale resources (e.g. EAGLE resolution of the Milky Way)
 - Algorithms? (Grav.+ Hydrodynamics, decomposed?)
- Experiments: LSST, Euclid, JWST, LOFAR, SKA...:
 - Data centric
 - Require substantial computing resources (pre- and postprocessing)
 - Comparable situation with LHC



Astro: scaling (HACC, N-body)





Astro: planets, stars, supernovae...

Examples:

- SN explosion
 - Resources: e.g. present PRACE call: 150 Mch
 - Understand the mechanism of the explosion
 - Predict neutrino and gravitational-wave signals
- Plantets, stellar systems
 - Turbulence of central importance in accretion disk
 - Resources: e.g. present PRACE call: 30 Mch
 Increased spatial resolution quickly reaches Exascale



Astro: Large Synoptic Survey Telescope (LSST)

- Dark matter
- Mapping the Solar System & Milky Way
- Data 1.3 PB/a
- Computing:
 - 100 Tflop/s initially
 - 250 Tflop/s long term





Astro: LOFAR

- 400 TB/d, reduced to 40 TB (post-proc.essing)
- Frequency parallel, mem. size limited, I/O bound
- Post-proc. can be parallelized,





Astro: SKA





Astro: SKA

Processing (O(100) Pflop/s = 4*Tianhe-2 = 71MW)

- CSP (central signal processor)
 - Real time data processing 12 TB/s (SKA1)
- SDP (science data processor)
 - Data parallelism (position, polarization, frequency, time)
 - Limited inter-process communication (some global ops required)
 - Data locality to minimize data movement
 - Accelerator design considered for compute nodes
 - Partial HPC network for global ops



Astro: Summary

- Present allocations
 - O(1) Bch allocated
 - Single projects with O(100) can generate O(1) PB data
 - Moving towards "realistic" parameters immediately calls for Exasale resources
- Experiments
 - Require substantial computational resources: 100 Pflop/s (SKA)
 - Large data sets, O(1) PB/a LSST, O(100) PB/a SKA





Plasma: Scaling (GYSELA)



S. Krieg | BDEC 2015



Plasma: core modelling requirements

Example: 5D gyrokinetic* simulation of plasma core transport in ITER, ¹/₄ torus; 1 ms run time

	Grid Size	Cores	core-h	Data (per run)
lon scale ¼ torus	300 x 10 ⁹	80k	22 x 10 ⁶	10 TB
Full torus	1.2 x 10 ¹²	160	100 x 10 ⁶	40 TB
Full confinement	1.2 x 10 ¹²	16M	100 x 10 ⁹	40 PB
Electron scale 1ms	10 ¹⁴	2B	360 x 10 ⁹	
Y. Sarazin et al, Nuclear Fusion 50 , 054004 (2010);				

http://www.fz-juelich.de/ias/jsc/EN/Expertise/High-Q-Club/Gysela/_node.html



Plasma: laser fusion energy – fast ignition



Challenging simulation with large computation and complex physics:

- Plasma density of 6X10⁴ n_c for laser fusion
 - cf: 0.01 n_c for electron accelerator; 100 n_c for ion accelerator. ($n_c = 10^{21} \text{ cm}^{-3}$)
- Multiple temporal scales: 3 fs for laser, 0.01 fs for plasma oscillation

Imholtz-Gemeinschat

Mitglied der Hel



Plasma: (hybrid) two-system Particle-in-Cell





Plasma: Fast ignition scenario req'nts

Example: hybrid PIC simulation of FI scheme; compressed target 300 gcm⁻³, box size 200 μm x 100 μm; 2 ps run time*

	Resolution	Size (df)	Cores	10 ⁶ core-h
2D	Δx=0.02 μm 50 particles/cell	50 M cells 2.5 B particles	16k	1
3D	Δx=0.02 μm 50 ppc	5x10 ⁹ cells 250x10 ⁹ particles	16M	100
3D hi-res	∆x=2 nm 100 ppc	5x10 ¹² cells 500x10 ¹² particles	16B	200 000

*Data from: Wang, Gibbon, Sheng & Li, Phys. Rev. Lett. **114**, 015001 (2015)



Plasma: Magnetic fusion (ITER)

Dedicated supercomputers for fusion community

2009-2013 HPC-FF (Jülich): 100TF

2013-2016 IFERC: 2PF / 72k cores = 0.03 GF/core

2017- Current Eurofusion call: 8PF /300k cores Capacity: 2620 M core-h from 2017



Plasma: HPC resources

	DoE INCITE Proposals (10 ⁶ core-h)	PRACE + national EU (10 ⁶ core-h)	EU facility- generated demand (10 ⁹ core-h)	Projected demand >2020 (PF-y)*
Tokamak physics	500	500	ITER >2018: 5	15
Laser fusion (fast ignition)	300	100	HIPER >2020 1	3
Petawatt laser sources	250	400	ELI (x3) 5	15
Space/astro physics	200	200	ESA missions 2	6
Total	1200	1200	13	40







HEP: LQCD spectrum



- Decays
- Fine-structure (iso-spin breaking)
- Multi-Nucleon systems: combinatorially and exponentially hard
- \rightarrow D, T, ³He, ⁴He, ⁶Li,...
- \rightarrow NN, NNN forces

Price-tag:

>20 + O(100) Pflop/s-yrs

mholtz-Gemeinschaft



HEP: LQCD, zero temperature

- a_µ
- Matrix elements
- **g**_A
- GPDs, DAs, TMDs
- Bag parameters
- CKM physics
- Transition form factors



Price tag:

O(100) Pflop/s-yrs

Mitglied der Helmholtz-Gemeinschaft

etc.



HEP: LQCD, zero temperature

- a_µ
- Matrix elements
- g_A
- GPDs, DAs, TMDs
- Bag parameters
- CKM physics
- Transition form factors



• etc.

04.02.2015













HEP: LQCD, finite T and μ



Finite temperature:

- Equation of state
- Critical point
- Freeze-out conditions
- Strangeness
- Fluctuations of con. charges
- Backgrnd. magnetic fields

>100 Pflop/s-yrs



HEP: LQCD

- Highly scalable codes: O(10^6) strong scaling demonstrated (e.g. "High-Q club, JSC")
- No "no-go" scaling hazards at this point
- Efficient solvers less scalable: used in (naively parallel analysis)





HEP: LQCD computers



→ QPACE 2 (SFB TRR 55), XeonPhi based

S. Krieg | BDEC 2015



HEP: LQFT, EFT

- Nuclei states
- Nuclear physics input for nucleosynthesis in supernovae
- Supersymmetric models
- Technicolor

Price tag:

 \approx 20 Pflop/s-yrs



Hoyle-state (red), PRL 106, 192501



HEP: LHC

- 2 stage Trigger system, 40M events/s → 1K events/s
- 4 Experiments: Atlas: 320MB/s, CMS: 220MB/s, ALICE: 100MB/s, LHCb: 50 MB/s (2010)
- Total data: 15 PB/yr (2010), 27PB/yr (2012)
- Data storage and analysis in tier-ed grid (WLCG):
 - Tier-0: @CERN, 5.5+17PB storage, 6000 CPUs
 - Tier-1: distributed, 11 centers, 1+10 PB storage, archive
 - Tier-2: distributed, 150 centers, analysis, no storage
 - Tier-3&4: WLCG access points



HEP: LHC

Data processing

- Streaming
- Serial execution (few cases of threading due to mem restrictions)
- Simulation code is being vectorized







HEP: LHC

Data processing

- Streaming
- Serial execution (few cases of threading due to mem restrictions)
- Simulation code is being vectorized







HEP: Summary

	present	2016-2020
LQCD, computing	ca. 5 Pflop-yrs	>500 Pflop-yrs
LQCD, data	O(200) TB	O(20) PB
LHC, computing	ca. 1 Pflop-yr	ca. 10 Pflop-yrs (extrapolated)
LHC, data	250 PB installed	ca. 600 PB (extrapolated)