Computing Infrastructure Projects for Particle Physics

PPAP community meeting 17/18 Sept 2012

Pete Clarke reporting from the Computing Advisory Panel

The Computing Advisory Panel (CAP)

CAP advises STFC on Computing "with a big C"

- HPC
- HTC
- Data
- Networking
- Data management and open data access policy
- Interactions with other bodies (e.g. PRACE)

Covers computing for all STFC science:

- PP Theory
- PP Experiment
- Nuclear Physics
- Astronomy
- Astrophysics
- ParticleAstrophysics
- Cosmology
- Solar Science
- STFC facilities

Members:

Pete Clarke (Chair)
Simon Hands
Roger Jones
Adrian Jenkins
Ralf Kaiser
Mike Watson
Jeremy Frey
Neal Skipper
+..facilities

CAP will therefore make a general submission covering all of these

Computing Projects relevant to PPAP

STFC funds two Computing Projects relevant to the PPAP sector:

GridPP:

High Throughput Computing (HTC) infrastructure for the LHC and other experiments.

DiRAC

High Performance Computing (HPC) for Lattice QCD

(note: DiRAC is more than PP)

Both of these are treated as projects in their own right and reviewed.

Both GridPP and DiRAC were requested to submit “project based” PR submissions

In last PR:

- *GridPP was ranked*
- *HPC operating costs were ranked*

GridPP

Thanks to Dave Britton for slides

GridPP provides Computing for all LHC experiments + other non-LHC experiments

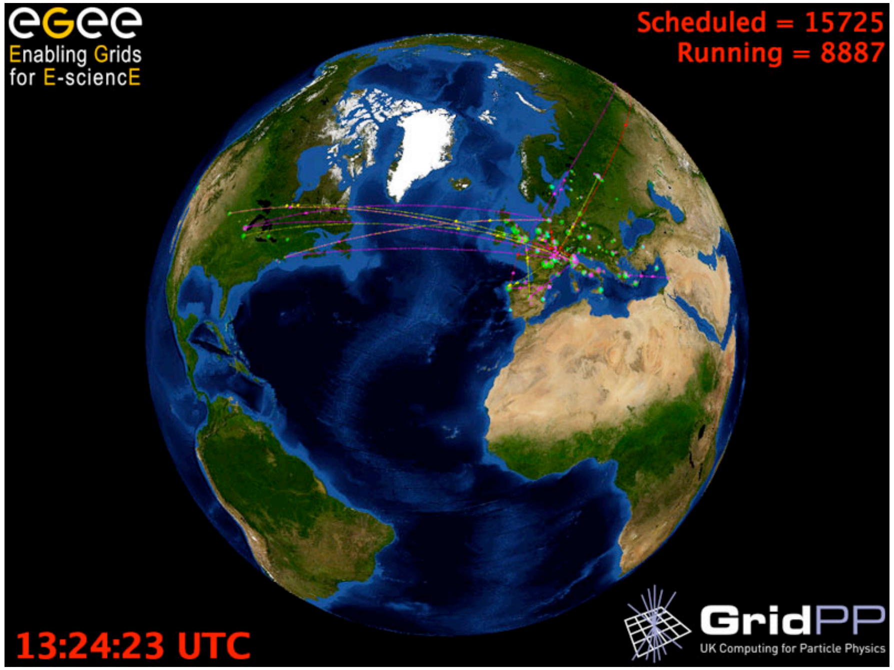

Part of WLCG (Worldwide LHC computing Grid)

Combined Resources (August 2012):

- 152 Sites in 36 Countries
- 325,000 logical CPUs
- 210 Petabytes of disk
- 180 Petabytes of tape

Architecture

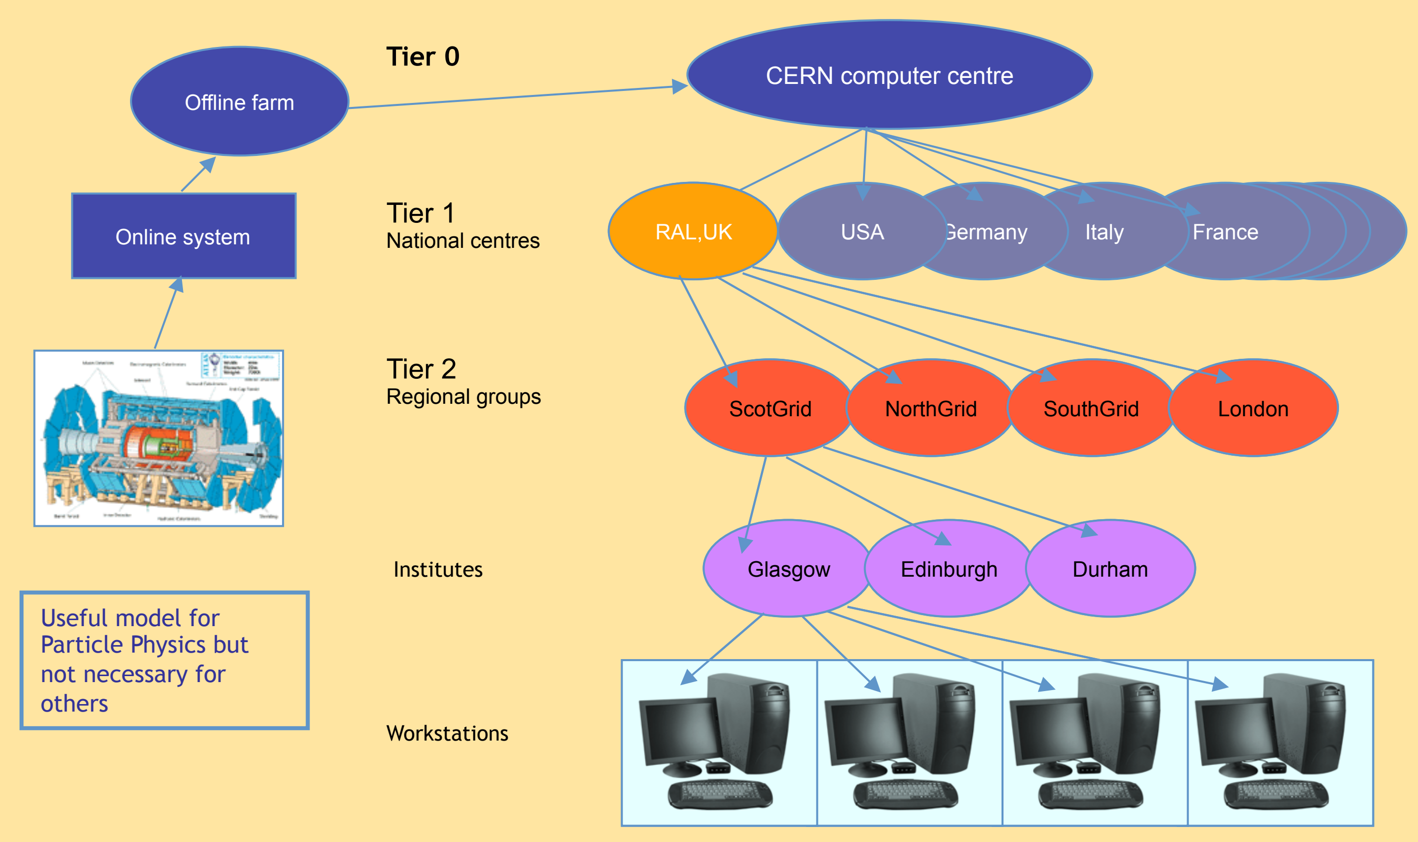

Useful model for Particle Physics but not necessary for others

GridPP: the UK Contribution

UK Resources (August 2012):

- 19 Sites
- 36,000 logical CPUs
- 21 Petabytes of disk
- 5+ Petabytes of tape
- 2 x 20 Gbit/s links to CERN

CPU contributions in 2012

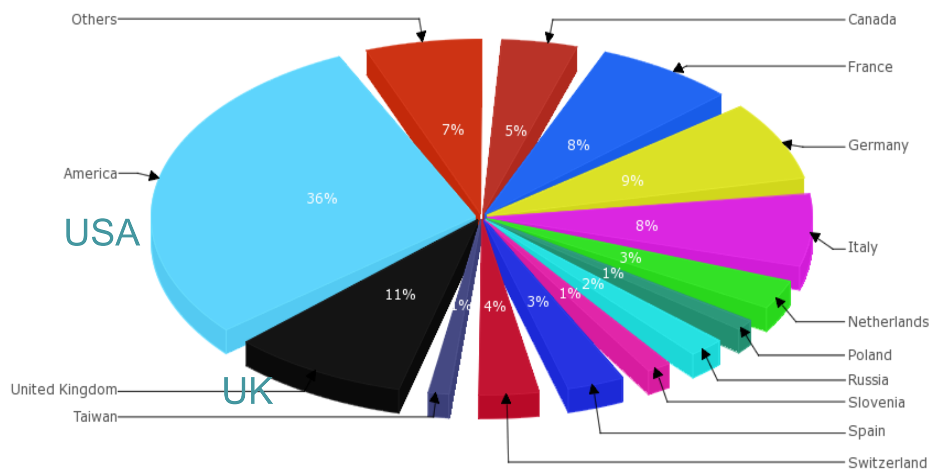

GridPP4 funded until ~ April 2015

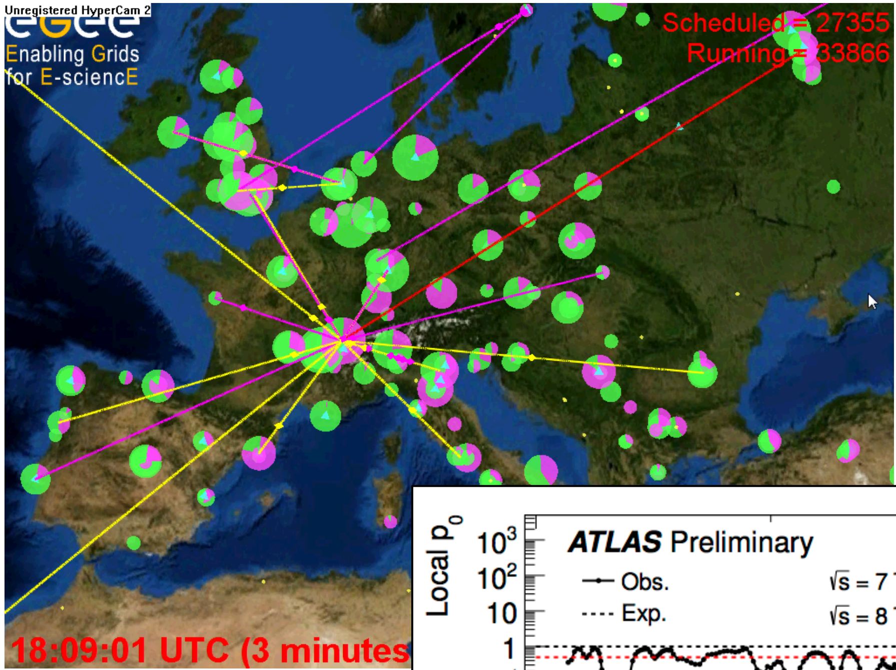

Results today only possible due to extraordinary performance of accelerators – experiments – Grid computing

Rolf Heuer

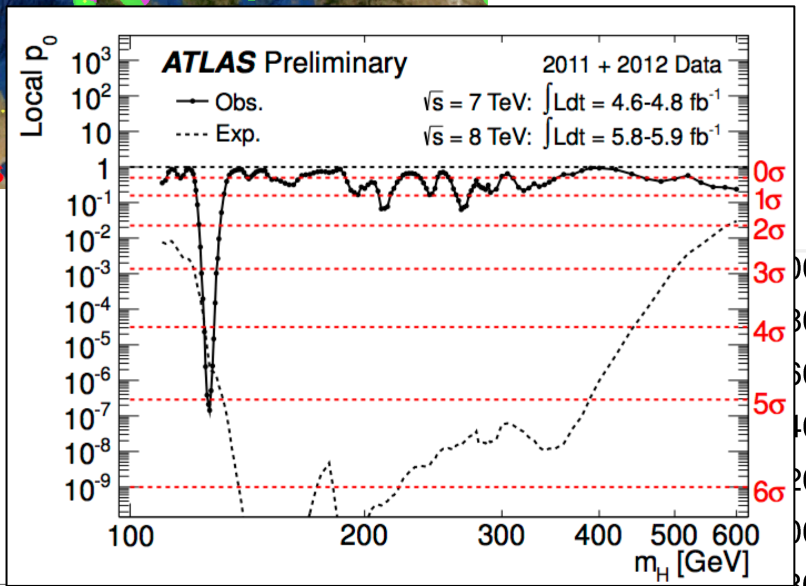

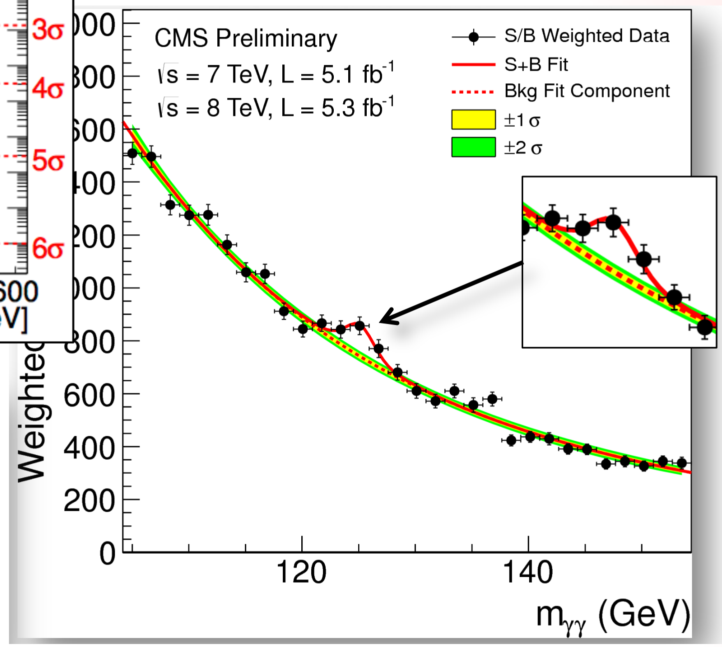

$B_{(d,s)} \rightarrow K^{-}\pi^{+}$

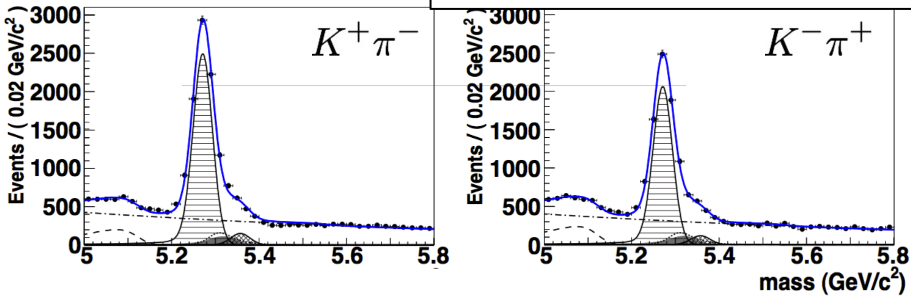

CAP recommendation:

- GridPP should be congratulated for the great success in handling the LHC data leading to recent LHC results.

- The GridPP computing infrastructure is essential for LHC exploitation

- GridPP should be supported in the future at the same priority as the science which it supports.

DiRAC

Thanks to Jeremy Yates and Christine Davies for slides

STFC DiRAC Facility

- Distributed HPC Services for UK Theory and Simulation activities in
  - Particle Physics,
  - Particle Astrophysics,
  - Nuclear Physics and
  - Theoretical Astrophysics
- Over last ~4 years DiRAC has coalesced HPC into a single facility - more effective and efficient use of funds and equipment
- 20,93 and 134 in tope 500
- Has now had 2 funding phases - now funded until 2015
- Peer reviewed resource allocation – open to all (including Nuclear Theory)

- UKQCD: 8 Universities ~ 50 people
- International collaborations : ETM, HPQCD, QCDSF, RBC-UKQCD, strongBSM

Five installations:

- Cambridge HPCS Service: Data Analytic Cluster - 10000 cores, 1PB Parallel File Store (PFS), Non Blocking Switch Architecture, 4GB RAM per core.
- Cambridge COSMOS SHARED MEMORY Service - 1856 cores, 14.8TB globally shared memory (8GB RAM per core),
- Leicester IT Services: Complexity Cluster - 4352 cores, 0.8PB PFS, Non Blocking Switch Architecture, 8GB RAM per core.
- Durham ICC Service: Data Centric Cluster - 6500 cores, 2PB usable PFS, High Performance IO and Interconnect, 2:1 Blocking Switch Architecture, 8GB RAM Per core.
- Edinburgh 6144 node BlueGene/Q - 98304 cores, 5D Torus Interconnect, High Performance IO and Interconnect.

System (supplier)	Tflops	Connectivity	RAM	PFS	Cost /£M
BG Q (IBM)	1260	5d Torus	16TB	1PB	6.0
SMP (SGI)	42	NUMA	16TB	146TB	1.8
Data Centric (OSF/IBM)	135	FDR10 (Qlogic)	56TB	2PB Usable	3.7
Data Analytic (DELL)	50% of 200Tflops	FDR (Mell)	19TB	2PB Usable	1.5
Complexity (HP)	90	FDR (Mell)	36TB	0.8PB	2.0

Lattice QCD hadron physics

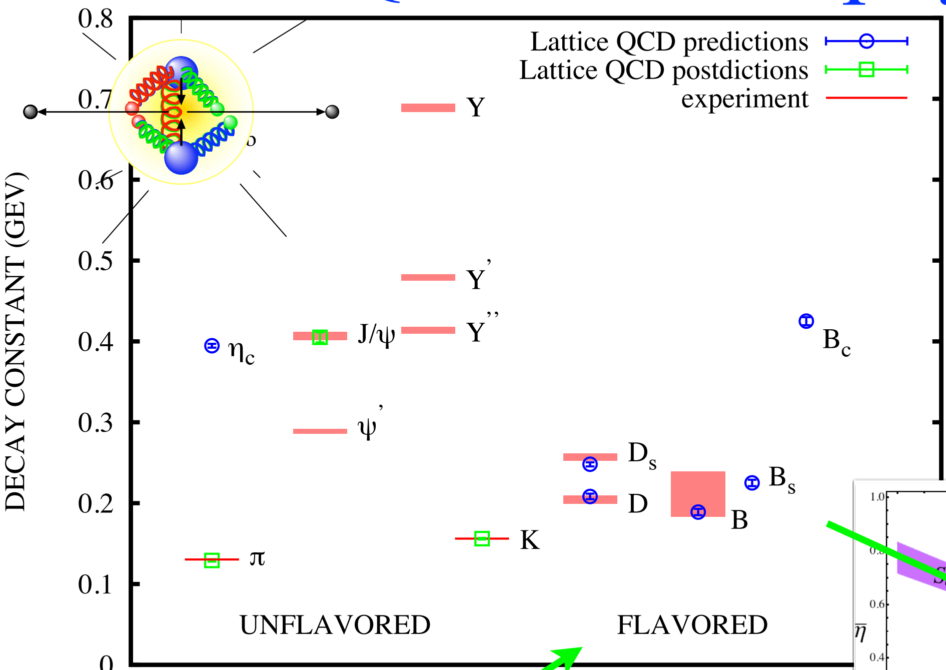

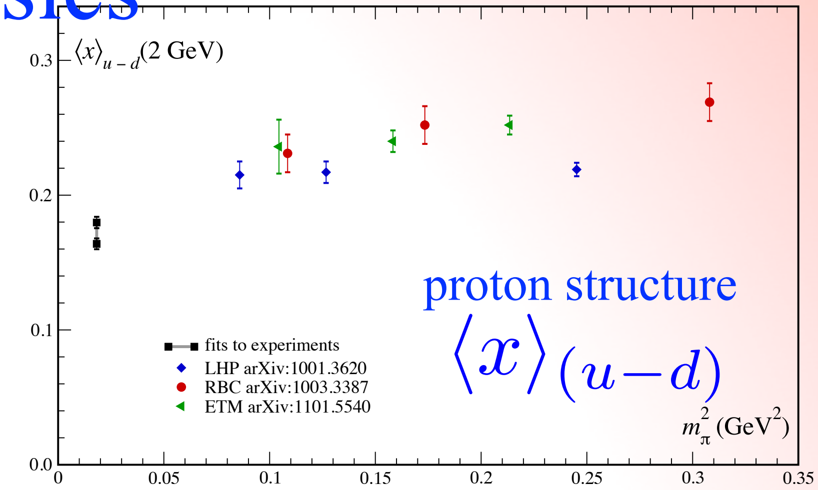

proton structure $\langle x \rangle_{(u-d)}$

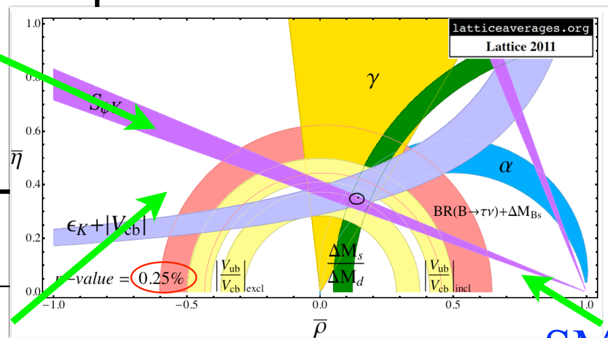

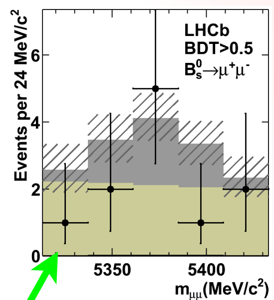

Precision electroweak MEs

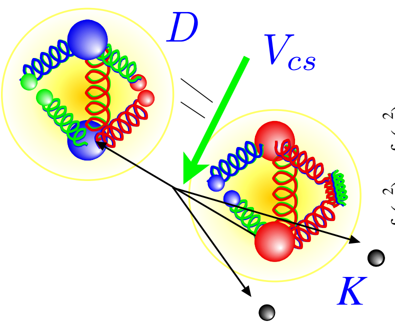

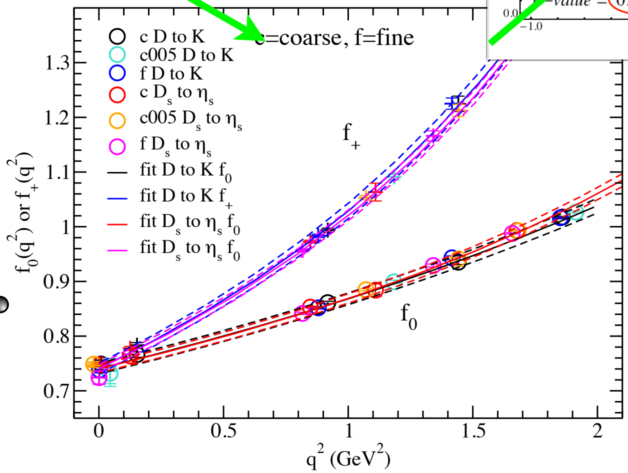

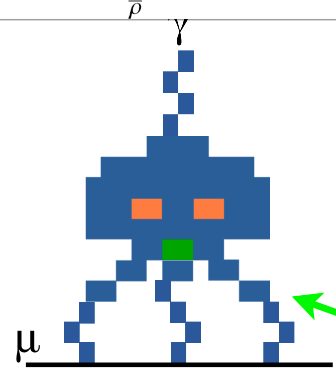

SM rates for hadronic EW processes need lattice QCD ....

muon g-2

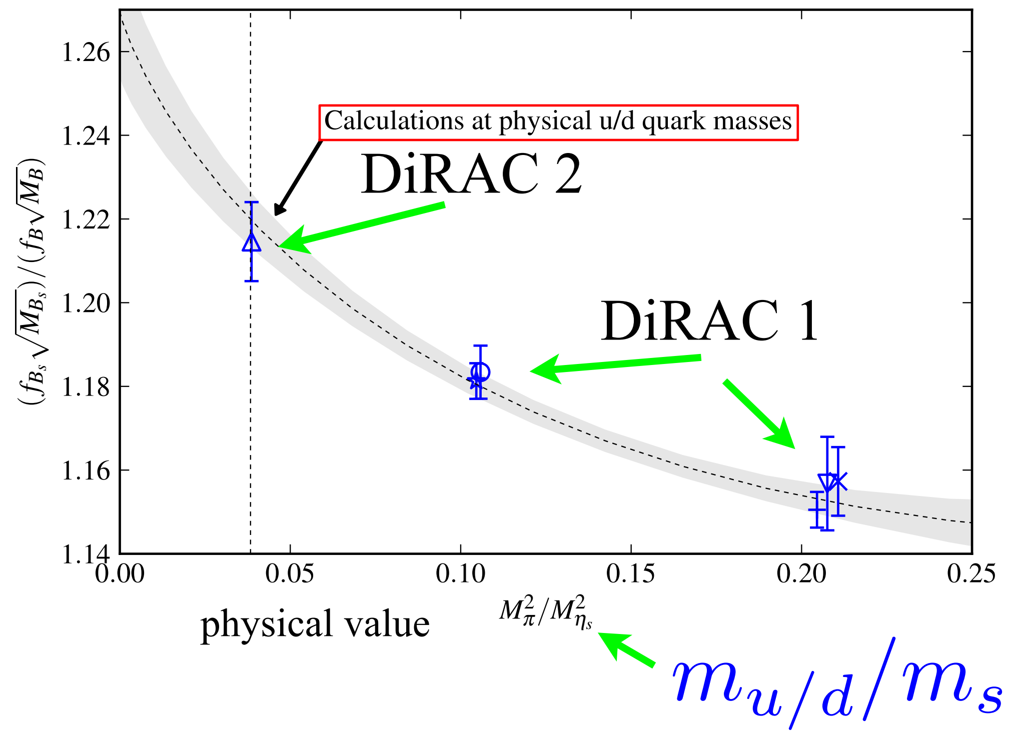

DiRAC 2 funded until 2015

Issue: Electricity only funded for 9 months !

CAP recommendation:

- DiRAC is congratulated for consolidating STFC HPC resources into an powerful and cost effective infrastructure well suited to the needs of its science base.
- The DiRAC consortium facilities are essential for all STFC HPC based science and should be supported in the future at the same priority as the science which it supports.

Other very important issues – I

Recurrent Costs:

- Recurrent costs of computing, particularly at institutes, have historically been dealt with an a bespoke and diverse way.
- Particularly for the HPC facilities this has made life difficult in the past.
- But the effect will pervade other areas sooner or later as Universities seek to establish more and more on MRF/SRF basis

A consistent message from CAP is that:

It is essential that STFC embraces the need to develop a consistent and holistic funding model which deals with both the capital and recurrent costs of computing.

Other very important issues – II

Data Management Policies:

- Policy for data management (preservation & open data access) is high in priority in Government, BIS, Research Councils.
- Policies exist in all Research Councils relating to both.
- In general Universities also have policies.

From the point of view of the science this has two consequences

- STFC supported projects will need to take some action at some point. At minimum have a policy. (LHC experiments are well advanced)
  ➔ CAP has been advising STFC in this sector.
- **More relevant today: These activities have a real resource request, and CAP has given the message that this should not be assumed to be "absorbed" in the project funding, nor that it can be done at minimal time cost by project physicists.**
  ➔ Curation for open access is a professional and non-trivial job.

Science Roadmap Challenge	DiRAC	GridPP	Astronomy Computing	Nuclear-Computi
A: How did the universe begin and how is it evolving?				
A:1. What is the physics of the early universe?	✔	✔	✔	
A:2. How did structure first form?	✔		✔	
A:3. What are the roles of dark matter and dark energy?	✔		✔	
A:4. When were the first stars, black holes and galaxies born?	✔		✔	
A:5. How do galaxies evolve?	✔		✔	
A:6. How are stars born and how do they evolve?	✔		✔i	✔
B: How do stars and planetary systems develop and is life unique to our planet?				
B:1. How common are planetary systems and is ours typical?	✔		✔	
B:2. How does the Sun influence the environment of the Earth and the rest of the Solar System?	✔			
B:3. Is there life elsewhere in the universe?				
C: What are the fundamental constituents and fabric of the universe and how do they interact?				
C:1. What are the fundamental particles?		✔		
C:2. What is the nature of space - time?		✔		
C:3. Is there a unified framework?	✔	✔		
C:4. What is the nature of dark matter?		✔	✔	
C:5. What is the nature of dark energy?		✔	✔	
C:6. What is the nature of nuclear and hadronic matter?	✔	✔		✔
C:7. What is the origin of the matter - antimatter asymmetry?	✔	✔		
D: How can we explore and understand the extremes of the universe?				
D:1. How do the laws of physics work when driven to the extremes?	✔	✔	✔	✔
D:2. How can high energy particles and gravitational waves tell us about the extreme universe?		✔	✔	
D:3. How do ultra-compact objects form, what is their nature and how does extreme gravity impact on their surroundings?	✔		✔	