



UNIVERSITY OF
BIRMINGHAM



From an IBM3090 to the Birmingham Environment for Academic Research in several difficult lessons ...

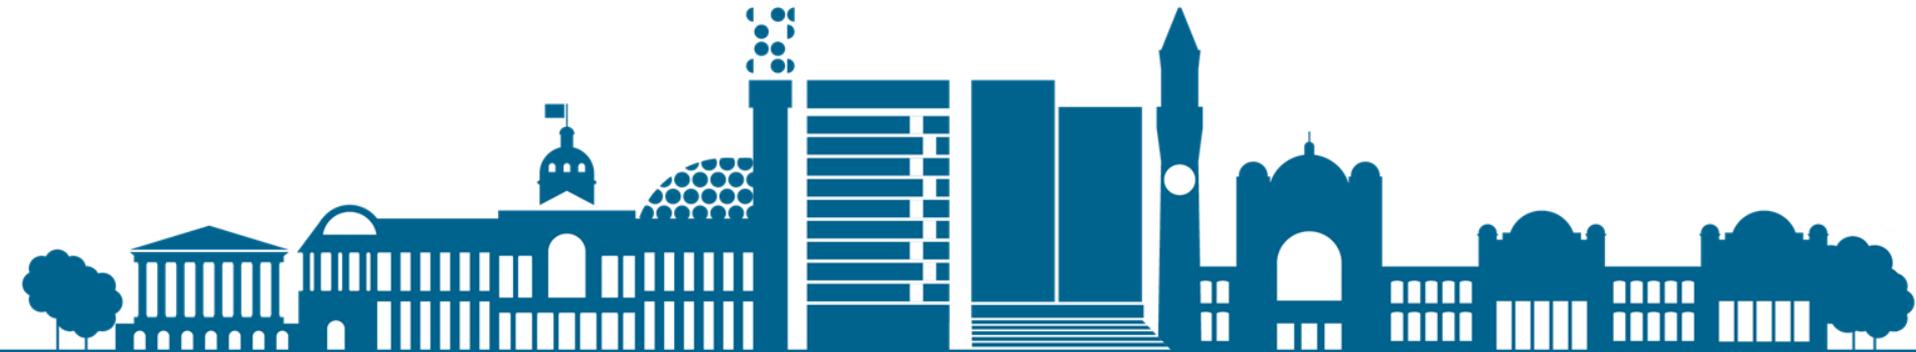
or

The BEAR Necessities

Paul Hatton

IT Services

The University of Birmingham, U.K.



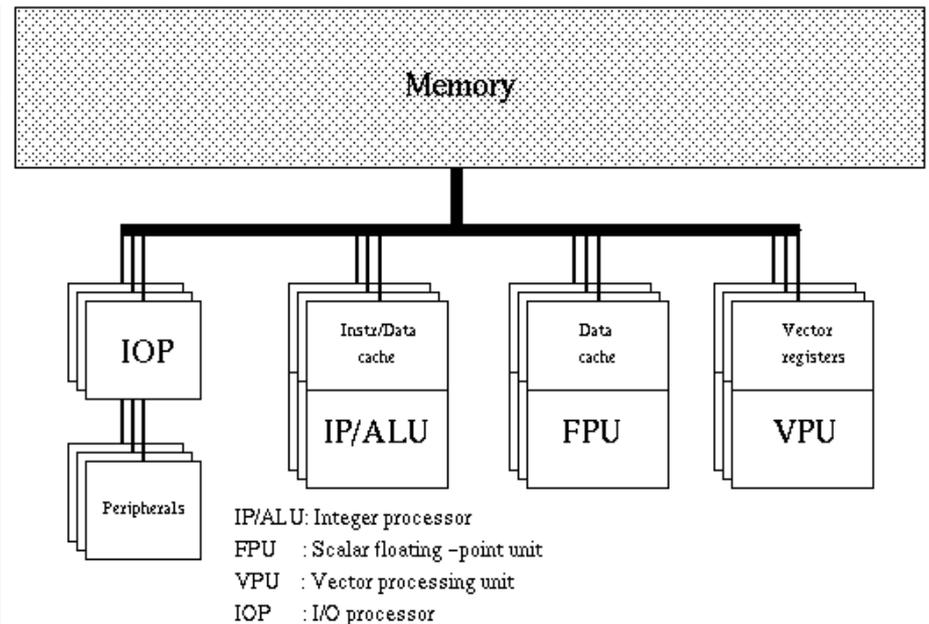
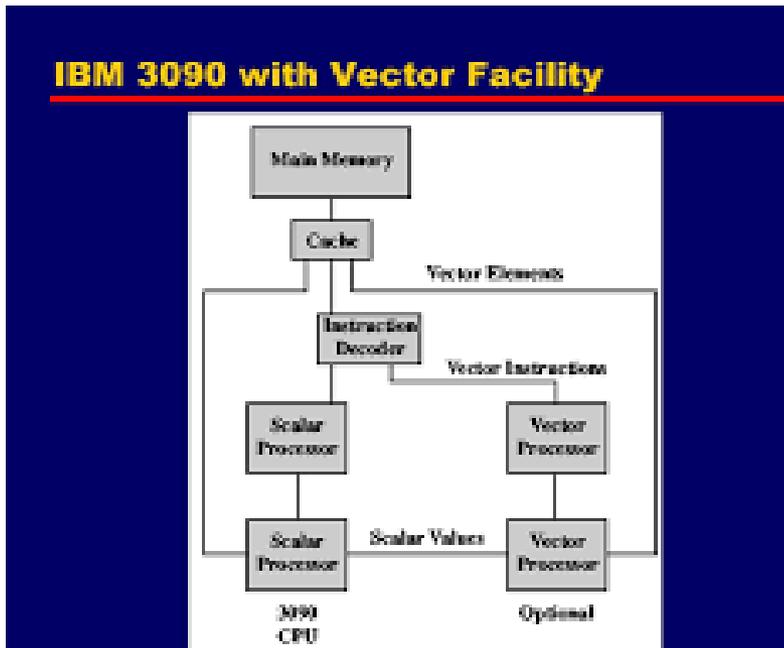
Nostalgia's not what it used to be

It all started with an IBM 3090 back in 1990



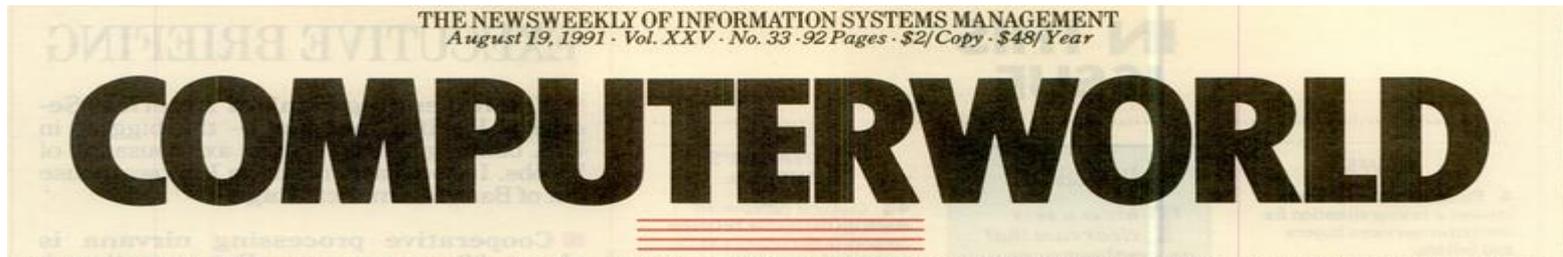
Nostalgia's not what it used to be (#2)

It all started with an IBM 3090 back in 1990
... with a vector facility



Nostalgia's not what it used to be (#3)

It all started with an IBM 3090 back in 1990
... with a vector facility ... and HSM ...



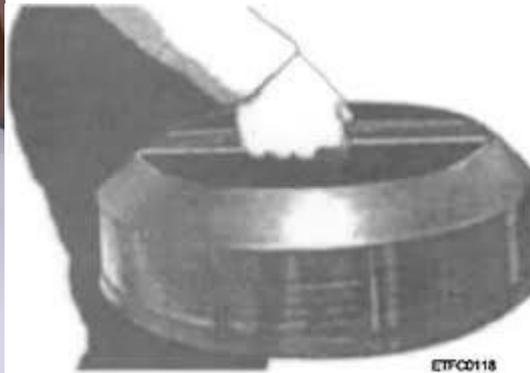
That was prior to 1987, when the company started using IBM's Hierarchical Storage Manager (**HSM**) for storage pooling. Since then, "we've improved disk utilization by about 15%," says Hansen, who oversees three job entry subsystem complexes of IBM 3090 Class 600 and higher machines.

In early 1990, Hansen boosted efficiency even further by adding IBM's System Managed Storage (SMS — see story page 51). With these products, Ford has been able to freeze — and even decrease — its magnetic disk capacity.

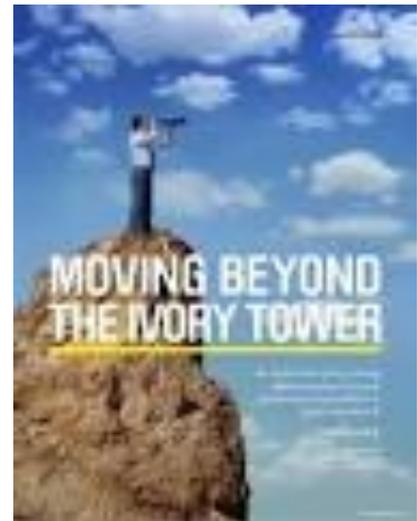


Nostalgia's not what it used to be (#4)

It all started with an IBM 3090 back in 1990
... with a vector facility ... and HSM ...
... and real disc drives ...



Times change



A little more nostalgia (1995)

- ❑ 2 x Compaq Alphaservers
- ❑ ISDU (Information Service Digital Unix) service
- ❑ VAX/VMS (we had to learn a new language)
- ❑ One for interactive work (isdugp, EV5 processors), one for batch Compute Intensive (isduci, upgraded to EV6 processors)
- ❑ We ended up as a bleeding-edge site for Compaq's cluster filestore
- ❑ isdugp service OK, isduci service unreliable due to flaky filestore
- ❑ no demand for parallel processing



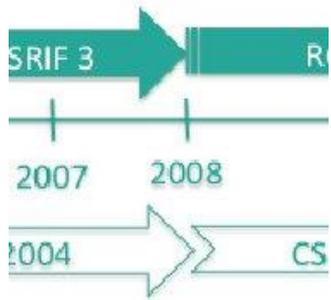


More recently (early-2003)



- ❑ Funded from SRIF-1
- ❑ HP Central Applications Service (CAPPS)
- ❑ 6 x HP J6700 servers each with 2 x PA-8700 processors at 750 MHz, max 16 GB/server
- ❑ HP/UX (we had to learn a new language)
- ❑ Funded as a general applications service, not specialist HPC
- ❑ Added researcher-funded (Archaeology) 1TB filestore, at great expense to the group – our first example of joint IT Services/researcher funding
- ❑ Unreliable filestore – basic NFS mounts to servers and IT Services SAN
- ❑ Users liked it when it worked but constant, and justified, complaints about poor availability



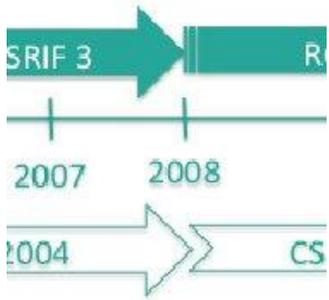


SRIF-3



- ❑ We don't need to discuss the limitations of the SRIF funding model (do we)?
- ❑ Our first foray into 'proper' HPC (2006)
- ❑ My first foray into the HPC-SIG
- ❑ Tried to be conservative, not bleeding-edge; we needed to provide a stable service
 - Selected dual-core AMD Opteron processors – quad-core processors were overdue and considered too great a risk by us
 - Selected GPFS as a supported cluster filesystem
- ❑ Scientific Linux – we had to learn a new dialect (from HP-UX), not a new language





SRIF-3



Didn't work out as hoped:

- ❑ No funding released by Birmingham for additional staff
- ❑ GPFS flaky – poor performance and reliability - and much buck-passing between integrators and IBM
- ❑ Integrators placed a specialist on site, who showed little interest in providing a stable service but loved tinkering
- ❑ Users liked it when it worked but constant, and justified, complaints about poor availability
- ❑ We couldn't carry on like this





a BEAR is conceived ...



We had to go back to basics since:

- ❑ We had a poor track record in providing a stable HPC or applications service
- ❑ We had little support from the majority of the researchers
- ❑ Many groups were running their own clusters with little incentive from the University to change
- ❑ The SRIF-3 cluster was heading towards end-of-life and there was little enthusiasm from the University to fund anything more than a nominal replacement; any additional staff posts were excluded from this nominal funding
- ❑ etc.

We had a novel idea of actually asking researchers what they needed, rather than us giving them what we thought they wanted

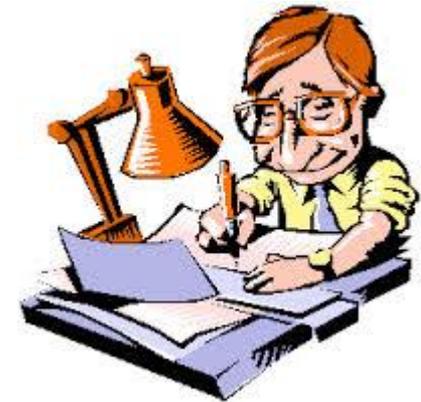
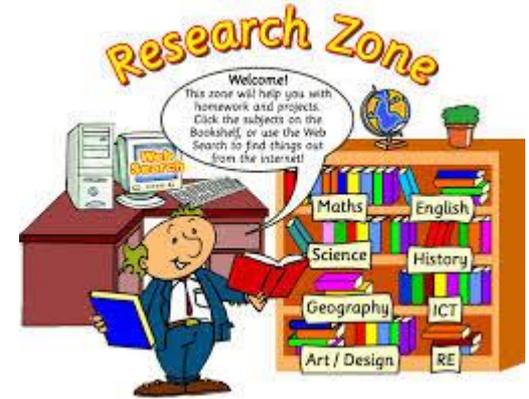


What do researchers do?

Some or all of:

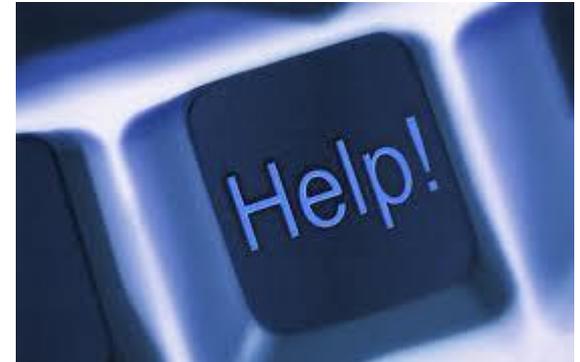
- ❑ Chase funding
- ❑ Learn the tools of the trade
- ❑ Collaborate
- ❑ Generate and consume data
- ❑ Compute
 - Other project-specific resources
- ❑ Interpret
- ❑ Disseminate

And, of course, the broader context



What are IT Services expected to do?

- Support a wide range of researchers with different expertise and requirements
- Lack of Linux expertise does not mean that research projects are not worthy of centrally provided resources



BEAR (Birmingham Environment for Academic Research)

- In American Indian mythology, the Bear is the font of all wisdom and the source of all knowledge
- The vision was to establish a Research Environment with multiple complimentary services





BEAR Mk 1 (a BEAR cub)



- University released (modest) 4 year recurrent University funding from 2012
- We went to Procurement for BEAR, not solely – or even primarily - for an HPC service
- The procurement made it clear that this was not a technology demonstrator but was for a stable and reliable service with a partner, not just a supplier
 - Partnership weighted more heavily than raw hardware in the procurement
- Set up as a partnership with OCF, IBM, Adaptive Computing, Mechdyne, IOCOM, Nice and NVIDIA
- We started to win back the doubters and build a reputation on campus as a trusted service provider





BEAR Mk 1 (a BEAR cub)



At this time we could offer:

- Linux-based HPC (on a good day)
- large (at the time) storage for HPC
 - but only local to HPC
 - data had to be transferred to/from HPC with sftp/WinSCP
- visualisation centre with active stereo and head tracking
- scalable collaborative conferencing
 - IOCOM Visimeet, based on Access Grid
 - think Skype on steroids
- **a vision**



What does a BlueBEAR cluster look like?





BEAR developments (a growing-up BEAR)



By 2016 the existing BlueBEAR cluster was heading towards end-of-life so another procurement:

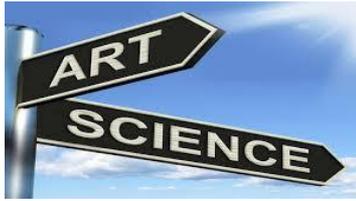
- ❑ With the new-found support from the research community the University released more recurrent funds
- ❑ We needed a long-term partnership to develop the existing services
- ❑ Went to procurement for an 8-year framework (4+2+2) years to give a long-term relationship
 - Getting this through procurement was quite a challenge
- ❑ Continuity with existing services was essential
- ❑ Resulted in a strengthening of our existing partnerships to give a firm basis for future developments



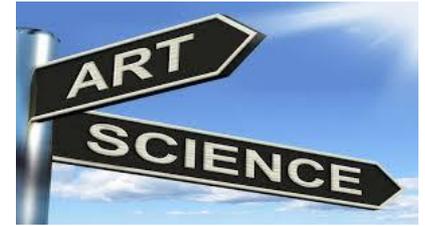
It seemed a good idea at the time

- Windows HPC
- Hadoop
- Batch render service
- And probably others





current BEAR services (1)



The University expects the services to serve a wide range of users, not just the usual HPC disciplines.

- BlueBEAR HPC
- BEARView which offers:
 - distributed and collaborative visualisation (NICE DCV)
 - collaborative conferencing (IOCOM)
 - high-end visualisation (IBM Visual and Spatial Technology Centre) with active stereo and head tracking
- MySQL database service
- hosting and maintaining servers for research groups
- training
- relationship with BEAR partners
- integration with regional and national centres





current BEAR services (2)



The University expects the services to serve a wide range of users, not just the usual HPC disciplines.

- Dropbox-like sync-and-share service, securely hosted in the campus datacentre
- Central storage
 - Research Data Store (RDS), common to all services for active projects
 - Research Data Archive, for long-term archive of research outcomes
- Research Data Network
 - Dedicated fibre network between large data generators (sequencers etc.) to the RDS
- BEAR Cloud (Private OpenStack deployment)
- supporting services, e.g. gitlab





Recent Developments



In late 2017/early 2018 the Research Computing group was awarded 6 new full-time permanent posts by the University, reflecting the confidence that the University and the research community has in our Research Computing Services

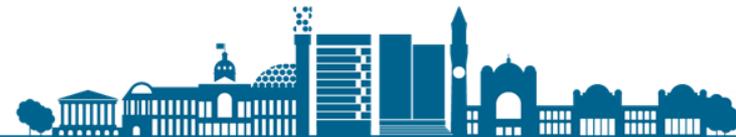
This was the largest growth in any comparable service in the UK academic sector

We have a long-term commitment to recurrent investment from the University to central Research Computing services - we do not have to bid annually for core funding, so long as we constantly make the case by highlighting our success and our value



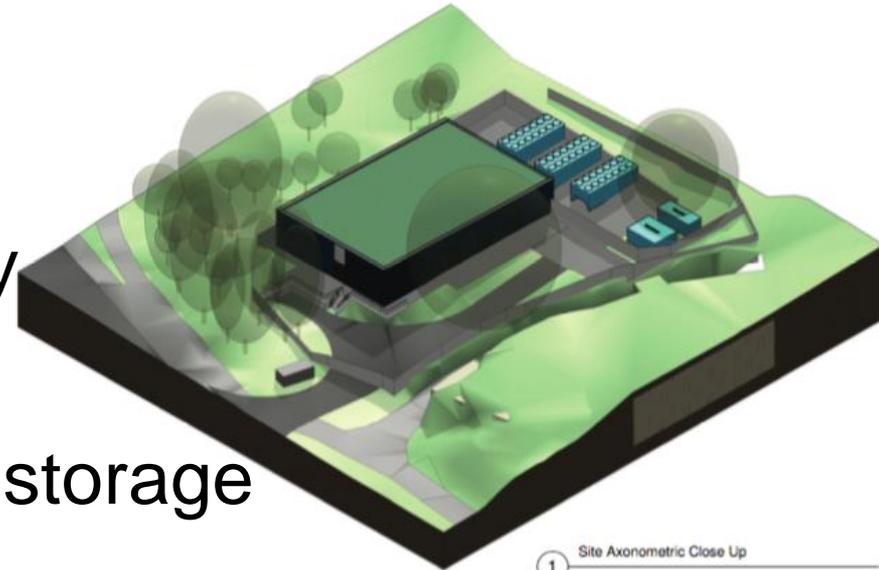
Team changes

- Added more staff 😊
- Architecture Infrastructure & Systems
 - 6 people + 1 student
- Research Engagement
 - 4 people
- Research Software
 - 4.5 people
 - + 1 student
- Planning student intern 18/19



New data centre

- Water cooling for all systems
 - (WCT/RDHX)
 - No air handling ...
- Dense compute
 - 1MW electrical capacity
 - 36 racks
 - Floor to support dense storage





Some research areas using BlueBEAR

Economics - Optimal voting rules for two - stage committees

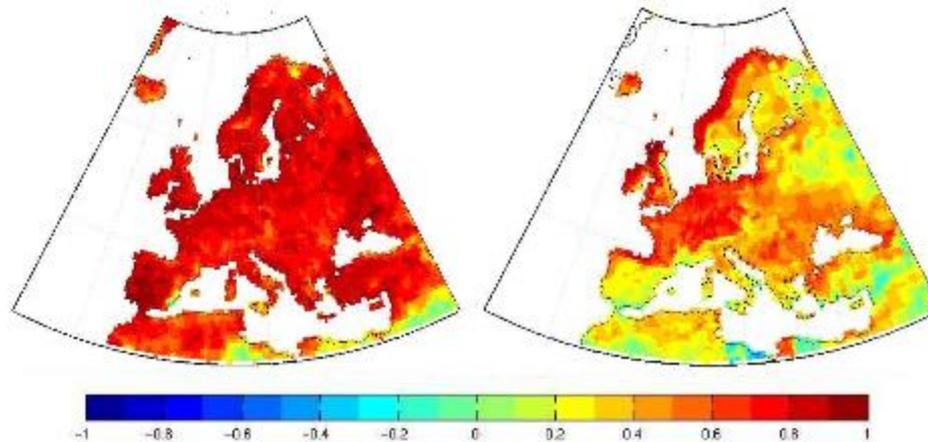


What voting rules should a committee adopt at the preliminary and final stages if it wishes to efficiently trade off its impatience against a desire to make the right decision?



Some research areas using BlueBEAR

GEES - Developing a correction for simulated rainfall from global climate change models.

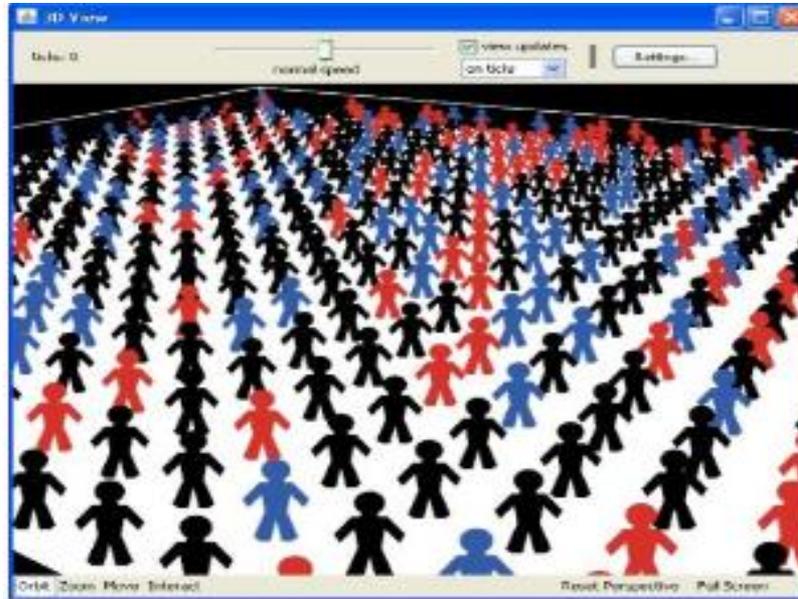


Our statistical corrections are shown to outperform some traditional down-scaling approaches and have great potential for application to future simulations.



Some research areas using BlueBEAR

Psychology - Simulating Social Stigmatisation

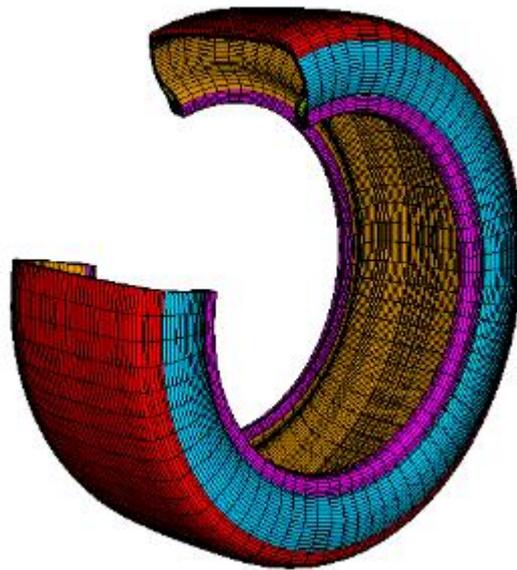


Existing psychological investigations tell us a lot about the group behaviours that can be observed in small convened groups, but there has been less work with larger groups and many of the emergent effects of small group findings have not been considered.



Some research areas using BlueBEAR

Mechanical Engineering - Developing Intelligent Tyre Based on Finite Element Analysis and Experiments

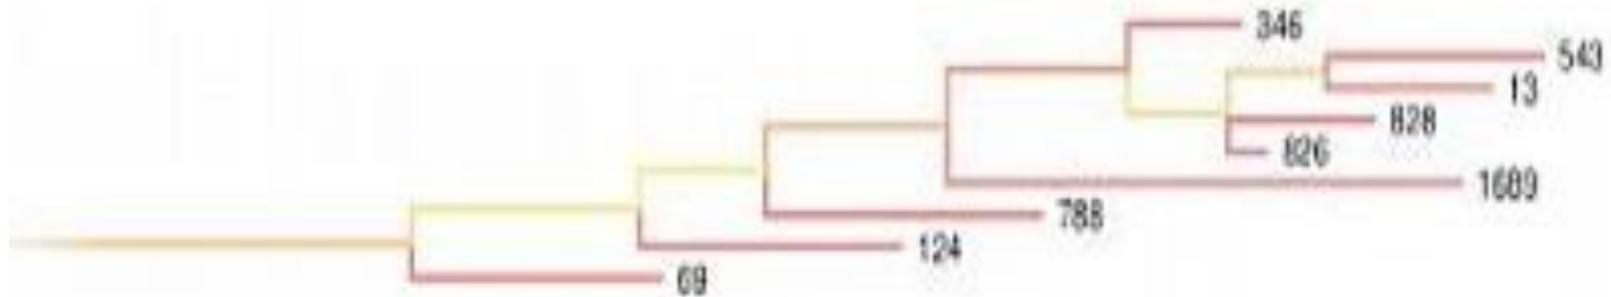


To develop a strain-based intelligent tyre system which can provide in time tyre information such as force and friction to improve vehicle safety, driving control and ride comfort, etc.



Some research areas using BlueBEAR

Discovering new family trees of biblical manuscripts



Until the arrival of the printing press, all documents had to be copied by hand – and every time a substantial document is copied by hand the copyist makes mistakes. For example the new manuscript might miss out lines, contain mis-spellings or even incorporate corrections to previous errors (or perceived errors) etc. These changes in the text are analogous to genetic mutations in species' DNA. Therefore the tools used in the field of phylogenetics to study the evolution of species can also be used to study the evolution of texts of a given work.



Some research areas using BlueBEAR

More at

<https://intranet.birmingham.ac.uk/bear/reports>



Future plans





Any questions?

