

Compute and data strategy meeting 8th October 2018

Present: Kenneth Heafield, Iain Rae, Steve Renals, Alastair Scobie, Ian Simpson, Amos Storkey

Purpose of group

The May 2018 meeting of the School's Computing Strategy Group agreed that :-

“... a working group should be formed to develop a strategy for resourcing the compute and data intensive needs of the School. This strategy would cover both teaching and research and consider :-

- an appropriate balance between physical and cloud provision
- the procurement of physical compute,
- the infrastructure required for supporting physical compute
- how we might resource the requirements of online courses
- a possible return to an SRF scheme

Projecting demand over the next two to five years would be an important first step before developing any strategy.”

The group is not tasked with solving any immediate problems with current compute clusters.

Current resources

Enumerating what we currently have :-

Over 500 GPUs with over 1,000,000 cores (as of Jul 2018) :-

- MSc teaching cluster
 - scutter – 12 nodes for hadoop
 - letha - 6 nodes providing total of 24 GPUs
- 200 GPU cluster (MLP)
- landonia - 25 nodes providing 200 GTX-1060 GPUs
- CDT
 - charles – 24 of nodes with typically 2 of Tesla or Titan X GPUs
 - james – 21 of 64 core nodes
 - two large memory (1TB) nodes
- Numerous research funded servers (both DICE and self-managed)

Growth in demand

As machine learning as a tool is being used by more and more informatics disciplines, GPUs are becoming the predominant compute platform.

The group believes that we do not have sufficient GPU capacity for our current, let alone future, needs.

Demand is expected to increase for the following areas :-

- Teaching (UG/PGT)
 - due to an increase in student numbers
 - * **ACTION: Ask Neil H for projected numbers of student over next 3-5 years**
 - due to an increase in use of machine learning in courses
 - * relatively few courses currently using GPUs, but we expect the number that do to grow substantially
 - * ditto for UG and PGT individual projects
 - * **ACTION: poll teaching staff asking about current and projected use of compute (traditional CPU, hadoop, GPUs, large memory) and storage for this year AND for next 3 years**

- Research (PGR and postdoc)
 - expect huge increase in PGR usage as machine learning spreads into other informatics disciplines
 - ditto for postdoc. (Although this use would be expected to be funded by research grants, the School still needs to consider how to meet this requirement)
 - **ACTION: poll PGR and research staff asking about current and projected use of compute and storage ...**

Many of the group are sceptical about the viability of distance learning. The group proposes that any distance learning courses make use of commercial cloud offerings and the costs of this be included in the course fees.

One member of the group suggested that we aim to follow Moore's law, doubling our GPU provision every 18 months. (Is that doubling the number of GPUs or the number of GPU cores? Doubling GPUs is impractical unless they substantially reduce in size.)

Problems

- Procurement
 - The majority of our GPU nodes have been integrated locally (server node purchased from Dell (using Select Server agreement) and GPU purchased from Insight (using IT accessories agreement)). We have taken this approach because there is currently no UoE level agreement from which we can buy GPU nodes. Unfortunately the resulting nodes are space inefficient (5U for 4 GPUs). This route obliges the School to carry the risk of incompatibilities between the server nodes and the GPUs. Recently we have been buying more space efficient GPU nodes (4U for 8 GPUs) from, for example, Scan Computers using the “three quotes” mechanism, but this is risky from a procurement perspective as we are likely to exceed the OJEU threshold. The “three quotes” mechanism is also time consuming.
 - There is a risk that nVidia could make it difficult for us to use commodity GPUs for research purposes (as it looked like they might have done early in 2018) either through restrictive licensing or by applying pressure to suppliers not to supply servers with commodity GPUs.
- Server rooms
 - Forum main B.02
 - * Space - there is currently space for 6 further racks. This could be increased by some re-structuring of existing racks.
 - * Power/cooling - the existing UPSes are close to capacity. Estates claim that there is at least 50% spare cooling capacity.
 - Forum self-managed B.Z14
 - * Space - this room is full with no scope for increasing capacity, short of evicting old space-inefficient kit
 - * Power/cooling - as for B.02
 - Forum self-managed B.01
 - * Space - there is currently 1.5 rack capacity free
 - * Power/cooling - power as for B.02. Cooling ??
 - Appleton Tower
 - * Space - with some minor rejigging, we have space for a further 2 racks
 - * Power/cooling - there should be sufficient spare capacity
 - Kings Buildings
 - * Space - with some minor rejigging, we possibly have space for a further rack
 - * Power/cooling - there should be sufficient spare capacity
- Computing staff effort
 - Procurement - a substantial increase in the level of School spend has resulted in delays in the various stages of procurement. This has been compounded by the use of the time consuming “three quotes” procurement route which has required the involvement of more senior staff than should be necessary. It is not clear whether this increase in spend was a one-off or the new base level.

- Commissioning - the rapid growth in the number of managed GPU nodes has led to a substantial workload increase on the Research and Teaching Unit. This has resulted in substantial delays in commissioning new nodes. These delays have sometimes led to some people choosing to self install and manage their GPU nodes. (Physically installing GPU nodes is labour intensive, requiring at least two people).
- In service care - responsiveness to problems has similarly suffered. The rapid growth has also meant that some configuration has not been scalable causing some inconsistent behaviours - work has been undertaken to retrospectively improve on this.
- Deploying services - the deployment of services such as slurm, gluster and hadoop has been under-resourced and, as a result, what have been effectively prototype services have been introduced as live services. Substantial effort has been expended over this summer to improve on these deployments, but much more remains to be done.
- No computing staff have been employed to support GPU computing (nor support the CDT resources). An additional computing officer is expected to be hired in early 2019.
- Cluster file system - the version of GPFS became unsupported. It would no longer compile under Scientific Linux 7 and our licenses were tied to old server hardware. This necessitated a change in technology - the work to research and deploy a replacement technology fell to the same staff member working on the other aspects of the GPU provision.
- Costs - some costs are hidden to us and this skews our decision making. For example, power and cooling for the School's server rooms are not charged to the School, but are charged for in the University's ACF data centre (or any commercial data centre).
 - **ACTION: what are the real costs to the University of our electricity and cooling use (plus infrastructure capital)**
 - **ACTION: how much would hosting in ACF / Pulsant cost?**
- Funding - some funding sources are capital limited (eg CDTs with max £10k capital) but permit significant spend on services.

The meeting agreed that the lack of computing staff resource is currently the single biggest problem..

ACTION: Consider what additional staff resources are required to adequately resource our existing GPU provision. Considering what additional staff resources are required for future GPU provision will depend on our deliberations of the nature of that provision.

Possible approaches

Continue in house

- We could take a more “hard-nosed” attitude to retiring old, often space inefficient, servers (and GPU nodes).
- We could host GPU nodes at the ACF, or even at a commercial data-centre such as Pulsant at the Gyle.
- We could standardise on GPU node hardware and bulk purchase. This should lower the costs and risks compared to our current procurement routes. It would also permit enforcement of higher density hardware which would help alleviate space issues.
- Integrating all our GPU resources into one infrastructure would facilitate greater resource sharing (teaching use and some research use is periodic) and would provide some cost efficiency benefits.
 - but we would need to have confidence that we can produce a very reliable infrastructure.
 We aim to demonstrate that this is possible using the MLP cluster.
- We could include GPUs in the teaching lab and PGR desktops. It was felt that these would only really be useful for code development, but not for real compute.
- We could consider organising an SRF (Small Research Facility) to cover GPU provision. An SRF would allow us to cover the cost of equipment refresh and would be more attractive to funding sources that are capital limited.
- It was noted that any solution must be capable of taking new model cards without necessitating a refresh of the host platform
- We could suggest to the University that it might discount or waive the costs of power and cooling

charges at ACF as part of match funding (eg for EPSRC grants)

- Ditto for GPU provision related to teaching on the basis of improving the student experience
- We could improve resource utilisation by providing a (web based?) service which gives accurate info on what GPUs exist and what they're used for, including self managed nodes.

Academic outsourcing

- IS
 - The EDDIE cluster has 32 GPU server nodes providing 44 Tesla K80s and 80 Titan Xs
 - IS took a very long time to deploy the EDDIE GPU provision.
 - IS are more inclined to purchase Tesla GPUs nodes, primarily because these are easier to procure.
 - As far as we are aware, nobody from Informatics has successfully run code on these nodes
 - **ACTION: Find out whether anybody has successfully used the IS nodes**
 - Members of the group have little confidence that the EDDIE team, as currently resourced, can deliver a viable GPU service
 - **ACTION: Discuss with IS their current GPU provision and plans for future expansion**
- EPCC
 - We are not aware of any GPU service provided by EPCC.
 - **ACTION: Discuss with EPCC their plans for GPU provision**
- JADE
 - The only national GPU service for academic use is the JADE service hosted by University of Oxford. This service consists of 22 DGX1s (providing 176 GPUs) and is reported to be already over subscribed.

Commercial outsourcing

Much of our GPU use could be satisfied using commercial cloud offerings from eg. Google Cloud, Amazon EC2, Microsoft Azure, IBM Cloud.

ACTION: cost out the various suppliers once we have realistic demand figures

Could we convince cloud suppliers to provide GPU time as a grant?

In practice, we are going to likely going to have to adopt a hybrid of the above offerings.

What are others doing

- **ACTION: investigate what Cambridge, Imperial and UCL are doing. Particularly wrt. teaching**
 - Cambridge have CSD3 Wilkes 2 cluster (360 of P100) funded by EPSRC with a variety of partner universities. We suspect not available for teaching use
- What do companies do? - we expect that they use spot instances, but these are considered of no use for teaching because of the unpredictability of their availability.

Last modified : 2018/11/22 15:07:20 Z