

Computation and Data Committee Report to the CASCA Board, May 2013

Current Committee membership:

James Wadsley (McMaster) (Chair)	Term ends: 31 June 2016
Hugo Martel (Laval)	Term ends: 31 June 2017
J. J. Kavelaars (HIA/NRC/CADC):	Term ends: 31 June 2014
Erik Rosolowsky (UBC/U Alberta)	Term ends: 31 June 2015

Jonathan Dursi has stepped down, as of fall 2013. Jonathan was very active and a great asset to this committee and to CASCA. Hugo Martel (Laval) has expressed his willingness to serve on the CDC committee in his place. Other current members have indicated that they remain willing to serve.

Computing in Canada

Historical detail on how High Performance Computing (HPC) has been funded in Canada, including the provincial consortia and the origin of Compute Canada was given in previous reports (see e.g. May 2013 CASCA CDC report).

Current Status of Compute Canada

Compute Canada (CC) has now been formally incorporated over a year, since Nov 2012. It is setting itself up to be the primary recipient of all national scale computing funds in Canada in place of the existing consortia. It currently has a mandate from CFI to manage most new equipment. In practice, CC has no centralized facilities so this currently means siting computers and staff within existing consortia. In future, CC may attempt to centralize systems and staff. After a hiccup involving hiring and firing an apparently highly qualified CEO (Bill Appelbe) early in 2013, CC has now has reinstated its interim CEO, Jill Kowalchuk (formerly of Cybera, a compute network organization in Alberta). It recently appointed an interim CTO (Chief Technical Officer), Jonathan Dursi (formerly of this committee of CASCA and seconded from Scinet Toronto), and CSO (Chief Scientific Officer), Dugan O'Neil (Physics, Simon Fraser), starting on 1st Jan 2014. CC has an aggressive and rather secretive board, chaired by Don Hathaway (no research or computing background). The board is mostly non-researchers. CC is set up as a not-for-profit with members who are a subset of Canada's universities and colleges (depending on the payment of a fee). CC also recently set-up an Advisory Council on Research (ACOR) last Fall. CASCA and NRC/HIA separately lobbied for representation. The member were chosen in an somewhat ad-hoc way (probably geographical). James Wadsley (McMaster), of this committee, is an Ontario representative and Robert Thacker (SMU) represents Atlantic Canada. Rob Thacker is currently chairing ACOR meetings (so far one by phone). Both were present at the 2013 AGM in Toronto. Detailed membership information and management personnel are listed on the website: <http://computecanada.ca>

CFI Funding Context

CFI has offered no new national High Performance Computing (HPC) funds since 2006. The last money was spend in 2009 or so, resulting in systems being installed in 2011. Most systems will be off warranty within a year and many will have to be shut down in 1-2 years. Some new equipment has appeared through CFI programs (LEF/NIF). Even though CC manages this hardware, it is dedicated to the use of specific groups. CFI provided new CC operating funding (though to 2015) which has enabled consortia to keep staff. A condition on this was a CC Strategic Plan being in place by March 2014. The proposed plan is high level. It includes no planning regarding new hardware or any

statements regarding researcher need. The draft strategic plan is here:

https://computecanada.ca/cc_files/news/CC_NationalConsultation_StrategicPlanDraftOutline.pdf

CC ran town hall meetings in December 2013. They are still accepting input by email:

consultation@computecanada.ca

CASCA first submitted a white paper (authored primarily by Jonathan Dursi) to this process in Jan 2013. With Jonathan Dursi as part of the CC management we feel our needs are well understood. However, the biggest problem is that CFI is not allowing applications for national level HPC infrastructure. Large scaling computing (new hardware costing over \$100,000) has been explicitly excluded from the new Innovation Fund call:

<http://www.innovation.ca/en/OurFunds/CFIFunds/InnovationFund>

On the other hand, CFI dictates that any (necessarily small) computing hardware funded must be managed by CC (which has frustrated CASCA members in the past). Major HPC re-investment has been deferred as part of a “cyberinfrastructure initiative” getting underway in 2014. CFI was clear that HPC would need to establish need in future and that no funds have been earmarked for major HPC infrastructure as part of this initiative. HPC clearly has an image problem in Canada.

Compute Canada Activity

The first CC AGM since incorporation was held in Oct 2013. The formal part was a 20 minute teleconference with very little discussed beyond the formal auditors report. Prior to the AGM, some members (VPR) had suggested CC would be “held to account”. However, though the member representatives on the phone expressed some frustration at a lack of communication, it was wrapped up very quickly with no discussion of governance or planning. After the AGM, there were presentations by Chad Gaffield (SSHRC president) and Robert Davidson (CFI VP Programs and Planning). Gaffield put forward a case for data stewardship and this has been reflected in recent tri-council reports. However, his focus is more on smaller social science datasets (Digital Humanities) and there was a major disconnect with the scale of data needs in the sciences. Robert Davidson expressed skepticism that the need was that great. Guillaume Bourque (McGill/ACOR) said his genomics lab alone would create data equal to CC's entire storage (~ 1 PByte), compromising investments in expensive sequencing equipment if new storage wasn't found. Davidson suggested Genomicists should pool resources and delete data. He said that 90% of researchers have stated they will be happy with just the computing they can get now for the next 7 years (based on a CFI survey apparently). This is ridiculously naive as current systems must be replaced in 1-2 years. Davidson talked about excellence as the driver: “CFI does not aim to float all boats”. He said he still didn't see a “Canadian reason” to invest more resources in HPC. Kowalchuk (CEO) remarked to the room that CC recent strategic planning was “not about getting money from CFI”. After Davidson left, James Wadsley asked the board why this was the case: “Isn't the strategic plan all about getting money from CFI?”. The board/CEO didn't give a direct answer. Astronomy is in a similar position to Genomics and other “big data” fields in needing new investment, particularly in longer term storage capacity. The attitudes of both CFI and CC are very disturbing. ACOR members used the meeting to push the board for lobbying – saying that new capital investment was critical to address aging equipment and new needs (e.g. Data). The board acknowledged the need but had no answers. In particular, CC has done nothing other than work with CFI. Under pressure, board chair Don Hathaway admitted no lobbying had occurred or was planned. He suggested that board members were individually influential people. For example, Fassi Kafkeye (Bombardier) might soon go on a trade mission and might therefore have a chance to talk to a minister.

At the current time, ACOR represents the primary research input into CC. The CSO, Dugan O'Neil, has yet to do anything and is relatively junior for such a role. ACOR was invited to the AGM to present researcher needs and advise CC on writing their strategic plan. ACOR members developed discipline-based PPT presentations of research computing needs but were given no opportunity to present them. As an advisory group, ACOR, is meant to be called on by the board, CEO or CSO and rather than be proactive. The members also appear to have a limited appetite for independent activity. There has been no interaction aside from the AGM. As noted above, Astronomy is fortunate to have Jonathan Dursi as part of the CC management. However, Jonathan has made it clear that he is first and foremost a CC employee from this point. At several town halls, issues of lack of transparency and local control were raised. However, the CC position seems to be that what has gone before was a failed experiment. In particular, CC is also drafting an "Optimization Plan" that may require some consortia to essentially shut down and a "Management Plan". ACOR has not been asked to contribute to this process. Current CC plans indicate ACOR will be consulted on the strategic plan in some way prior to the formal adoption. There is no formal interaction between CC and the provinces or the CC membership (VPR's at institutions) that might provide other channels for input.

A key concern is that CC is moving slowly and being entirely reactive – only responding to CFI requests and failing to pursue other funding. As of writing, there is no active lobbying to counter the negative view CFI (or government) has of national HPC. There is some activity in individual provinces. James Wadsley is researcher representative on the SHARCNET board of VPRs and was thus able to find out about Ontario activity. Ontario is the only province with more than one consortium and is working on incorporating a Compute Ontario which would provide a single voice that is a counterweight to Compute Canada. Ontario's Ministries are well aware of research computing needs and the government is considering a proposal to invest in HPC infrastructure without CFI/CC. This reflects considerable ill-feeling that some provinces currently have toward CC. Essentially, provinces are asked to pay 40-60% of the costs and yet CC wants to centralize decision making. VPRs in Ontario have become quite concerned. They are working with the province but have also proposed that professional organizations in the science and engineering could play an important role, particularly if they help make governments aware of how serious the situation has become.

The 2005 HPC Long Range Plan was modeled on the Astronomy LRP and had a high impact (albeit on a different government). It led to the 2006 CFI NPF program. That plan was written by researchers. No current planning or other efforts by CC seem to fill this niche. Given the negative perception that there is of Compute Canada (or any existing HPC organization) by the national government, input that was perceived as independent could be very important going forward. CAUT offered to help facilitate lobbying with organizations such as CASCA involved. A concern there is that CAUT may be viewed negatively by the federal government which may undermine the message. CFI has already begun some consultations with select researchers in the lead up to this cyber-infrastructure initiative. However, based on their current misconceptions, it isn't clear that their consultation processes work. There are some indications that the selection of researchers is somewhat biased toward the digital humanities with no traditional big HPC users involved.

Impacts on Astronomy

Theoretical astronomers are commonly running parallel simulations and their students are producing theses that rely on them in areas like star and galaxy formation, planetary dynamics, CMB and numerical relativity. As HPC systems get shutdown or become congested, the only possible response will be less ambitious research. This will make it harder for our students to get noticed and to continue in the field. Until recently, simulators have typically been able to apply for time and get CPU time

allocations within an order of magnitude of competitors in Europe, Asia and the US. However, competing groups are now pulling away. In Germany and the US they have now run astrophysical simulations for a single project that represent of order 20% of Canada's entire HPC capacity. As this trend worsens, Canada will be unable to make the simple claim of being competitive.

Observational Astronomy has been surging forward into the “big data” regime that is considered new by many other disciplines. CANFAR has established world-class capabilities in managing big data sets but is not being matched by required infrastructure. It should be noted that local consortia (e.g. SHARCNET) are well aware of this evolution in the need for storage vs. cycles and but don't have much total capacity because their systems were planned over 7 years ago.. Thus overall, CC has not progressed on creating a long-term storage solution for large data sets. The CANFAR project receives a yearly allocation and is requesting more storage each year. The project is honest about the fact that this is long term storage. However, CC doesn't have a plan of how to make that happen. In the most recent allocation they expressed their concern that the storage situation was ‘not scaleable’ and that has us worried about future storage requests. Data hosting is one way Canada can participate in large projects without needing direct investment but without new HPC investment we are missing these opportunities (e.g. PANSTARRS).

Despite the lack of support for new astronomical investment projects in Canada as a whole, there are still several large data projects that continuing or ramping up. CANFAR was extended, as was the CyberSKA, both through the CANARIE NEP programs. The ALMA Development Program has offered to support further development of remote visualization and archive tools at the several FTE-year level. ADASS will be hosted at the University of Calgary the week of Oct. 6, 2014. Thus Canada possesses a lot of expertise and would be able to do a lot more with appropriate infrastructure investments. If we could provide firm numbers regarding the level of need and how specific projects would be compromised it would help support a lobbying effort. As of writing, it still isn't clear how to get such a lobbying effort going. A major problem is the lack of leadership within Compute Canada, which would be the natural organization to coordinate this effort.