Large-Parallel Supercomputer Simulations – Frontiers in Canadian Research

R. Fernández[1] (UAlberta), F. Herwig (UVic), A. Babul (UVic), W. East (Perimeter), C. Groth (UToronto), N. Ivanova (UAlberta), P. Kushner (UToronto), L. Lehner (Perimeter), C. Matzner (UToronto), P. Myers (UAlberta), J. Navarro (UVic), P. Navrátil (TRIUMF), D. Peltier (UToronto), H. Rein (UToronto), D. Siegel (Perimeter / UGuelph), R. Thacker (SMU), J. Wadsley (McMaster)

Executive summary: Supercomputers like Niagara or its much more potent international siblings are an indispensable tool to understand complex, multi-scale, and multi-physics real-world problems. Modern supercomputers allow predictive, first-principles simulations of reality by computing the relevant fundamental laws of nature. The authors of this white paper lead research programs involving such simulations in aerospace, astrophysics, atmospheric/ocean/climate physics, and nuclear physics. A similar methodology is also applied in areas such as material science and engineering. These large-scale simulations are now often the only way to create theory-based predictions that can be compared with modern, detailed observations, and are therefore an essential component of knowledge creation in these research areas.

In order to perform these simulations the **first** important prerequisite is **a dedicated large, national supercomputer consisting of the order of ten thousand homogeneous nodes,** tightly integrated with fast interconnect, and **maintained and operated to support large-parallel computations, each involving tens to hundreds of thousands of processor cores.**[2] **Second, a human resources funding infrastructure to build and continuously advance** – in close collaboration and union with the scientists – **the software code instruments that enable these simulations**. Third, such simulations create very big and valuable data sets for which a dedicated cyberinfrastructure is required to enable efficient, collaborative data analysis and exploration. While large-scale simulations share the latter requirement, for example, with physical experiments and data-driven approaches, the first two – dedicated large-parallel supercomputers and support for the creation of sophisticated software instruments – are specific and of particular importance to this approach to scientific innovation.

With the *Niagara* supercomputer, the Canadian scientific community has had for the first time access to an internationally competitive simulation facility, enabling numerous new lines of research in multiple fields. It has been possible, for example, to carry out high-resolution simulations of global ocean dynamics with up to 48 000 cores, which are able to resolve bottom ocean topology and the creation of waves and nutrient mixing. It is of utmost importance to build on this success and not to lose ground with respect to our competitors in this key emerging methodology of scientific discovery. **We recommend an ambitious renewal and expansion path for the large-parallel simulation capability in the Canadian DRI ecosystem, to build on existing progress, and to expand the associated specialised human resources required for scientific discovery.**

1. Canada's supercomputing science at the international research frontier

The *Niagara* supercomputer allows distributed memory calculations of up to 80 640 cores, with an estimated peak performance of 4 peta flops[3] to tackle demanding problems in science and engineering. Due to its unique design (e.g., fast interconnect) and its size, it is the only Canadian supercomputer that handles jobs larger than 1 024 simultaneous cores in the regular scheduling mode. Computations are supported by 12 Pb of spinning disc storage (*project* and *scratch*), a burst buffer with 256 Tb of NVMe-based fast storage, and an associated tape storage facility (*HPSS*) with current capacity of more than 30 Pb.

With this type of computing facility, we can address the challenges that are common among our disciplines, such as a *large dynamic range* in spatial or temporal scales, where processes operating at one scale modify the behavior of a system on a very different scale, and where a high spatial or temporal resolution – and

[1]Designated Contact Person (rafernan@ualberta.ca)

[2]For this type of high-performance computing, accelerator-based architectures relying heavily on GPUs are presently not serving the needs of this community, see Section 2 and 3 for more details.

[3]`https://www.scinethpc.ca/niagara/`

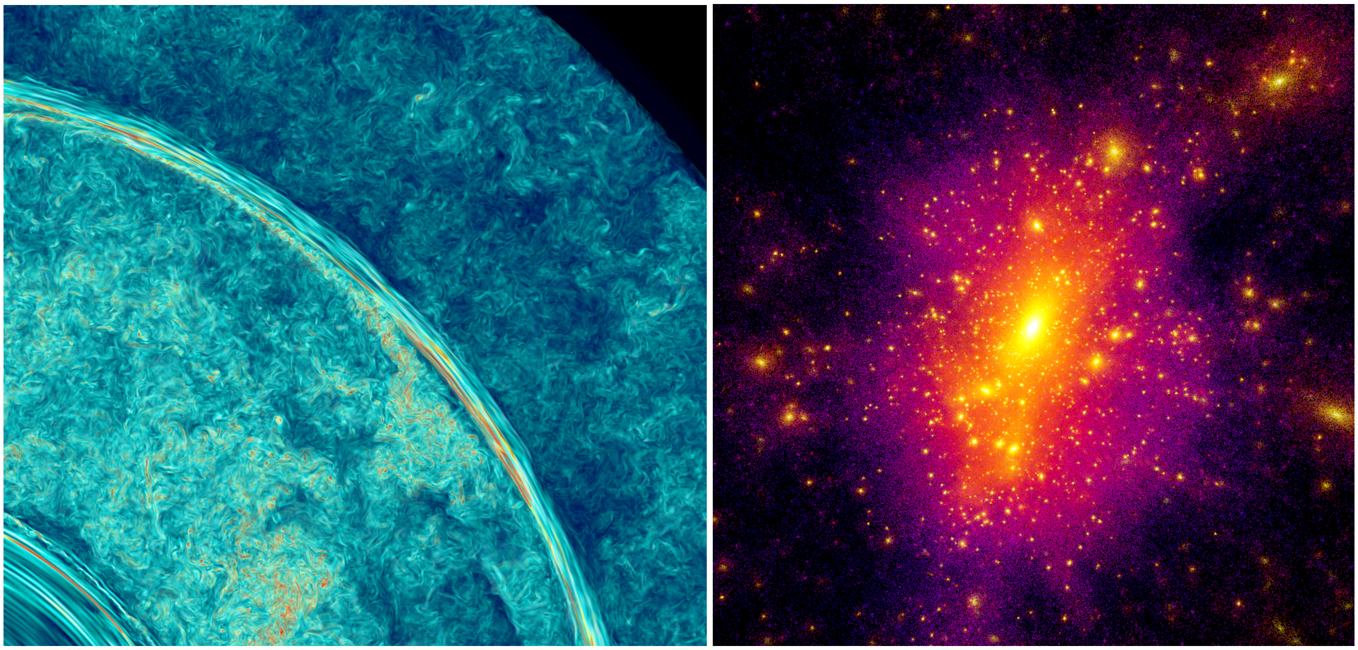

Figure 1. *Left*: Simulation of the turbulent hydrodynamics inside a massive star of the first generation that formed after the Big Bang in the nascent Universe, performed in 2019 on NSF *Frontera* using 400 000 cores by Herwig and collaborators using the `PPMstar` code. Shown is the vorticity of less than a quarter of the central plane of the full-sphere 3-dimensional simulation on a 3840^3 grid. *Right*: Dark matter distribution for an ultra- high resolution super-cluster from the Romulus suite of simulations. Romulus runs have been carried out using up to 100 000 cores on the NSF *Blue Waters* system with the massively-parallel tree+SPH code `ChanGa` (T. R. Quinn, A. Babul, Romulus collaboration).

therefore a large number of interrelated and concurrent calculations – is required to capture this behavior. *Global climate simulations, which address a problem that is mission-critical for all of humanity, provide an excellent example*[4]. Our simulations of the global ocean and atmosphere circulation and its climate evolution require capturing the dynamics on large (~ 1000 km, Earth size) and intermediate scales (100 km, such as ocean eddies and atmospheric mesoscale systems). Inclusion of smaller-scale ($1-10$ km) phenomena, such as that occurring in the Labrador Sea and atmospheric convective circulations, can fundamentally alter the large-scale circulation dynamics. Focusing on the ocean, small-scale phenomena have important links to the ocean's biogeochemistry and thus life in the ocean, and play a role in setting deep water formation and air-sea gas exchange. Ocean models capturing the smallest scales require resolutions of the order of 1 km, and typical numerical experiments on Niagara employ 5 000 to 20 000 compute cores for weeks to months to simulate behavior over multiple years.

Other common challenges are *complexity*, *non-locality*, and *dimensionality.* In astrophysics, some of the most pressing theoretical problems combine the challenge of large dynamic range in space and time with multiple physical processes simultaneously acting and influencing each other on different time scales, some of which are spatially local (e.g., fluid dynamics, micro physics) and some far-reaching (e.g., gravity, radiation). Key Canadian astrophysics strengths include galaxy formation and evolution, mixing and nucleosynthesis in stars, explosive transients including black hole and/or neutron star collisions and their multi-messenger signals, and the formation and dynamics of stellar and planetary systems.

A particular challenge involves simulating real-world problems in three dimensions (3D) to accurately project the correct dynamics, e.g., when turbulence is important. For 3D problems, the most practical and commonly adopted approaches dictate that the computational cost scales with the fourth power of the number of grid cells per spatial dimension. For example, a turbulence simulation on a 2000^3 grid typically

[4] See Climate Science 2050: Advancing Knowledge on Climate Science

requires 256 times as much computing resources compared to a simulation on a 500^3 grid. Typically, a factor of two in grid resolution is used to judge sensitivity to numerical effects, and a factor of 4 is required to provide a reliable measure of *numerical convergence*, a key element of simulation result verification. Factors of 256 in computational effort can only be realized within acceptable time if massively-parallel computing approaches are adopted. Often we find indications of new processes during parameter studies at low- or medium resolution. However, to be certain about these results, numerical convergence is needed, requiring 100× more computing resources than are regularly needed by most teams. Simulation results without verification have limited impact in the scientific community.

Simulations such as those shown in Fig. 1 would be possible on a next-generation *Niagara* successor that we recommend. For example, the $3840^3 = 56.6$ billion cells of the stellar hydrodynamics simulation shown in the left panel of Fig. 1 are advanced by 18 time steps per second by having 400 000 cores operate in concert using the NSF *Frontera* system. Only with such large grids can the small-scale turbulent instabilities in convection zones be accurately described. These small-scale turbulent instabilities constitute a critical mixing mechanism that determines how the elements were made in the first massive stars.

Candian researchers also carry out *ab initio* nuclear theory simulations that predict properties of atomic nuclei from basic interactions among nucleons. These calculations rely on a numerical solution of the quantum many-body problem, which requires large computational resources. Results enable the understanding of exotic nuclei investigated at rare isotope facilities like ISAC/ARIEL at TRIUMF, understanding nuclear reactions important for future energy generation and for astrophysics, and the testing of fundamental symmetries in nuclear processes. Computations demonstrate very good scaling beyond 6000 parallel tasks. This particular application also benefits from GPU accelerators.

2. International context and collaborative research

With the installation of *Niagara*, Canada has made progress in catching up with the large-parallel resources available to researchers in other industrialized countries. However, it is now essential to plan further expansions and make large-parallel computing a permanent and growing fixture in the Canadian DRI ecosystem. In several countries, notably the USA, large-scale scientific computing has a long tradition related to its importance for national security. Some of the largest systems in the world can be found at US national labs, such as *Summit* at Oak Ridge National Laboratory and *Sierra* at Lawrence Livermore National Laboratory, which combined are more than 60× as powerful as *Niagara*. Although some of these US Department of Energy (DOE) facilities are open to academic research use (such as *Summit*), these two machines are based on GPU accelerator hybrid architectures, which are not suitable at this time for most multi-physics simulations described in this white paper. For more than a decade, however, this type of computing in the US has not been limited to national labs. With the world-leading *Kraken* (2008), *Blue Waters* (2012) and *Frontera* (2019) clusters, the NSF has shown a long-term commitment to academic research enabled by supercomputing. *Frontera* was in 2019 the number five on the Top-500 list and 9× more powerful than *Niagara*, and is now – after *Niagara's* upgrade – still 6.8× more powerful. With *Niagara*, Canada has started to follow the path of providing researchers with internationally-competitive resources approaching what is needed, although with a decade of delay. **The most important recommendation of this white paper is that NDRIO continues to remain on this track of regularly and predictably developing and advancing Canada's large-parallel supercomputing resources,** so that Canadian scientists can continue catching up with their international peers. For example, a *Niagara* successor in place by 2023 with an increase of computing power by five to eight times would enable truly competitive high resolution simulations of the coupled climate system.

As of November 2020, the *Niagara* cluster ranks 82nd in the Top-500 list[5]. Several machines on this list are not available to academic or public-interest research, however, but instead belong for example to large oil and gas companies. Nevertheless, just like the US, many other industrialized countries prioritize dedicated large-parallel supercomputers for basic and applied public research, with one or more systems

[5] `https://www.top500.org/site/50206/`

more powerful than *Niagara*. This list includes[6] Japan, China, Germany, Italy, Saudi Arabia, Switzerland, France, UK, South Korea, Australia, Taiwan, UAE, Russia, Spain, Finland, and India. To highlight a few examples, six-times smaller Finland and Sweden provide a national research supercomputer that is respectively 2× and 0.8× as powerful as *Niagara*. Germany, with just over twice the population of Canada, features eight national machines ahead of *Niagara* in the Top-500 list, which are in combination 40× more powerful – and all of these are dedicated to academic or public research.

It is essential that Canadian researchers do not always turn to their international collaborators when large simulations are required. Due to their complexity, many of the most important large-parallel research programs can only be carried out by specialized international collaborations. It is a simple reality that those participants of collaborative efforts who can contribute meaningful computing resources to a project shape the scientific agenda and have the most impact.

Niagara enables Canadian researchers to participate in advanced international research in the first place. For example, in the JPL/Caltech - UMichigan - UToronto (PI Peltier) joint development of direct simulations of the ocean internal wave spectrum, *Niagara* is used to downscale the results from a global ocean simulation, initially run on the NASA *Pleiades* system, to force a regional version of the same model to attempt to reproduce the spectrum of ocean internal waves. Published results are based on 20 000−40 000-core *Niagara* runs. *Niagara* has also enabled significant grants, such as the $3.5M Canadian Sea Ice and Snow Evolution Network (NSERC CCAR CanSISE, 2013-2019, PI Kushner), which partnered seven Canadian universities and Environment and Climate Change Canada, relying on *Niagara* to generate multi-century numerical experiments that elucidated the global atmosphere-ocean response to Arctic and Antarctic sea ice loss. Also, because of the ability to run stellar hydrodynamic simulations for nuclear astrophysics on *Niagara*, UVic (PI Herwig) was able to participate in the NSF Physics Frontier Center Joint Institute for Nuclear Astrophysics and generate $0.85M in additional research funding from outside Canada.

Canada declared a climate emergency in 2009 and there is an urgent need to reduce greenhouse gas emissions (GHGe) due to the vast and intensifying human and environmental costs of climate change. Canadian high-performance computing facilities do somewhat better than other jurisdictions (e.g., Australia[7]) when it comes to its GHGe. However, we recommend to set an ambitious target date for when Canadian high-performance computing facilities will become net zero carbon. The computing resources used by researchers should not only be measured in core-years, but also in the amount of emissions generated, to raise awareness that large-scale simulations generate significant GHGe. We recommend that NDIRO set and publicize expectations around such GHGe reporting.

3. Broader DRI context and impact on personnel recruitment, benefits to Canada

Large-parallel simulations generate large amounts of data ($> 10-100$Tb per simulation). Rapid, short-term access of this data is required for analysis in post-processing, responding to peer review, and/or sharing data with collaborators. The data sets produced with the largest simulations can often answer more scientific questions by other researchers than those that originally motivated the simulation. One way to enhance the impact of this science is to provide public access to the data. But even the reduced and compressed data sets are too large for download, and specialized tools in combination with the capability to apply the tools to the data are needed. Furthermore, this actionable, analytic data access must be shareable at different levels beyond the Canadian community, to project the flexible and evolving international collaborative structures described in Section 2.[8]

Dealing with these technical challenges requires a broad spectrum of skills, and represents a rich DRI training environment with many meaningful roles for young researchers at all levels. *Niagara's* technical

[6] Excluding countries with single machines obviously associated with the private sector.

[7] Stevens, A. R. H et al., 2020, Nature Astronomy, 4, 843.

[8] An example of such a gateway that combines several of these requirements is the *Cyberhubs* platform (Herwig, F. et al., 2018, ApJS, 236, 2) currently deployed on the Arbutus cloud (`https://astrohub.uvic.ca`). It provides analytic, research-grade access to data from large-scale astrophysics simulations generated on *Niagara*, and was developed in part with cyber-infrastructure support from Canarie. Developing platforms like this requires support typically not available through Tri-council funding.

support team maintains and develops a specialized knowledge base derived from operating Canada's national supercomputer. *Niagara* and this ecosystem enables the training of graduate students and other highly-qualified personnel in the skills required to carry out large-parallel calculations, such as parallel programming, understanding computational efficiency, designing and operating sophisticated scheduling strategies, post-processing of large data sets, and effective long-term archiving of these large data sets. These transferable skills are of direct application to information technology, including to fields like data science or advanced visualization, which are in high demand by industry.

The complex software instruments (computer programs) for large-scale simulations are developed over decades. For example, the FLASH code used by some authors of this white paper simulates high-energy density phenomena such as supernova explosions[9]. Development started in 1998 as part of DOE's Advanced Simulation and Computing Initiative (ASCI) at the University of Chicago, and has since consumed 120 person-years for development and maintenance. To operate, adapt, and develop these codes requires a very specialized training at the interface of the science domain, the hardware architecture, and the appropriate programming models. A long-term plan for Canada to succeed in this area requires flexible human resources support for groups to develop and operate their domain-specific code instruments as well as post-simulation cyber-infrastucture and data pipeline. Support needs to go into creating groups and networks involving dedicated staff with expertise in parallel computing software and data analysis, working alongside researchers in an environment that attracts and trains the next generation of supercomputing researchers. The creation, adaption, and optimization of the code instruments that turn compute cycles into science are currently not sufficiently supported. Beyond the operation and maintenance of the machine itself, funding for staff and training of students working in the research groups is required. Therefore, our **second important recommendation is that funding instruments are created that provide flexible support for creating, optimizing and adapting the specific software and big data ecosystems required by large-parallel simulations and their shared use by the broader community.**

Code development and modification costs are especially relevant when considering new accelerator-based, hybrid architectures. As outlined in Section 2, other countries have historically developed a broad community capable of porting codes to new experimental architectures. Until such dedicated and continued support mechanisms for code instruments are firmly established in Canada, future large-parallel procurement should select a conservative, homogeneous CPU design based on established technologies and a *user-supplied basket of benchmark codes to ensure broad usability*. This is the philosophy that NSF has followed with its past machines, and its latest computer *Frontera* is indeed a quite similar design compared to *Niagara*. Cognizant of the large (and in many cases for academic users insurmountable) development costs demanded by more experimental, hybrid architectures (such as *Summit*), NSF appears to leave such designs to the US national labs.

Finally, a top-tier facility for large parallel calculations is a key prerequisite in recruitment, training, and retention of graduate students, post-doctoral researchers and faculty. Daniel Siegel, one of the co-authors of this white paper, is a recent faculty hire at University of Guelph/Perimeter Institute said "I would not have moved to Canada if there wasn't a large computer". Several authors were able to hire CITA National Fellows from premier international institutions (Columbia, Max Planck) explicitly because of access to *Niagara*. Before the arrival of *Niagara*, two international faculty recruitment attempts at the Canada Research Chair level failed at one of the co-authoring institutions in large part because of the lack of a suitable supercomputer for this type of research to be successful.

Final recommendation: This research community shares a common methodology of large-scale supercomputer simulations, representing a broad and deep resource of knowledge and experience in a key area of DRI that is increasingly recognized in all leading, industrial countries as a critical capability to solve applied, national security, and basic research challenges. **We recommend that the large-parallel supercomputing community is represented in the NDRIO Researcher Council, and that decisions on operation of current and design of future facilities will seek our input and involvement.**

[9] Dubey, A. et al. 2019, Int. Jour. Hi-Perf. Comp. Appl., 33, 322.