

NDRIO White Paper: Envisioning Digital Research Infrastructure for the Simons Observatory

Adam D. Hincks^{*1}, Simone Aiola^{†2}, J. Richard Bond³, Erminia Calabrese^{‡4}, Andrei Frolov⁵, José Tomás Gálvez Gherzi³, Renée Hložek^{1,6}, Matthew Johnson^{7,8}, Mathew S. Madhavacheril⁸, Moritz Münchmeyer⁸, Lyman A. Page⁹, Jonathan Sievers¹⁰, Suzanne T. Staggs^{§9}, and Alexander Van Engelen¹¹

¹Department of Astronomy & Astrophysics, University of Toronto, 50 St George St, Toronto ON M5S 3H4, Canada

²Center for Computational Astrophysics, Flatiron Institute, 162 5th Avenue, New York, NY 10010 USA

³Canadian Institute for Theoretical Astrophysics, University of Toronto, 60 St. George Street, Toronto ON M5S 3H8, Canada

⁴School of Physics and Astronomy, Cardiff University, The Parade, Cardiff CF24 3AA, UK

⁵Physics Department, Simon Fraser University, 8888 University Drive, Burnaby BC V5A 1S6, Canada

⁶Dunlap Institute, University of Toronto, 50 St George Street, Toronto ON M5S 3H4, Canada

⁷Department of Physics, York University, 4700 Keele Street, Toronto ON M3J 1P3, Canada

⁸Perimeter Institute for Theoretical Physics, 31 Caroline Street N, Waterloo ON N2L 2Y5, Canada

⁹Joseph Henry Laboratories of Physics, Jadwin Hall, Princeton University, Princeton, NJ 08544, USA

¹⁰Physics Department, McGill University, 3600 Rue University, Montréal QC H3A 2T8, Canada

¹¹School of Earth and Space Exploration, Arizona State University, Tempe, AZ 85287, USA

December 15, 2020

Abstract

Observations of the cosmic microwave background (CMB) are an incredibly fertile source of information for studying the origins and evolution of the Universe. Canadian digital research infrastructure (DRI) has played a key role in reducing ever-larger quantities of raw data into maps of the CMB suitable for scientific analysis, as exemplified by the many scientific results produced by the Atacama Cosmology Telescope (ACT) over the past decade. The Simons Observatory (SO), due to start observing in 2023, will be able to measure the CMB with about an order of magnitude more sensitivity than ACT and other current telescopes. In this White Paper we outline how Canadian DRI under the New Digital Research Infrastructure Organization (NDRIO) could build upon the legacy of ACT and play a pivotal role in processing SO data, helping to produce data products that will be central to the cosmology community for years to come. We include estimates of DRI resources required for this work to indicate what kind of advanced research computing (ARC) would best support an SO-like project. Finally, we comment on how ARC allocations could be structured for large collaborations like SO and propose a research data management (RDM) system that makes public data releases available not only for download but also for direct analysis on Canadian DRI.

1 Introduction

For more than a decade, Canadian digital research infrastructure (DRI) has been central to advances in cosmology through the study of the cosmic microwave background (CMB), the oldest light in the Universe. Compute Canada (CC) has supported the analysis and processing of microwave data from experiments such as the *Planck* satellite, the South Pole Telescope (SPT), the SPIDER B-mode experiment, the Atacama B-Mode Search (ABS) and the Atacama

*Corresponding author. Email: adam.hincks@utoronto.ca

†Simons Observatory Data and Pipeline Project Leader

‡Simons Observatory Theory and Analysis Committee Chair

§Simons Observatory Spokesperson




Figure 1: The Atacama Cosmology Telescope. The Simons Observatory will be at the same site. (Photo: A. Hincks.)

Cosmology Telescope (ACT). The challenges of analysing data sets from these diverse experiments differ. As an example, we will take as a baseline the computational needs of ACT, which has been a significant user of Canadian DRI since 2009, and then discuss the transition to its next-generation successor, the Simons Observatory (SO). Future telescopes such as the Cerro Chajnantor Atacama Telescope-prime (CCAT-prime), the *LiteBIRD* satellite and the CMB Stage Four (CMB-S4) telescopes (c.f., Hlozek et al., 2019) will have analogous computing needs, though with varying data volumes. The recent Canadian Astronomical Society Long Range Plan for the coming decade listed Canadian participation in a large ground-based CMB survey like CMB-S4 or SO, as well as the space-based *LiteBIRD*, as priorities (Barmby et al., 2020). DRI support will be critical to realising these goals.

ACT is a 6 m telescope in northern Chile that observes the CMB at high resolution over wide areas in multiple frequencies (Fig. 1; Swetz et al., 2011; Thornton et al., 2016; Henderson et al., 2016). Raw data have been reduced to maps of the sky using resources managed by Toronto’s SciNet Consortium, starting with the GPC cluster in 2009 and continuing on its successor Niagara. These maps have enabled almost all of the ~ 70 scientific papers produced by the collaboration, including groundbreaking results such as the first detection of the power spectrum of gravitational lensing of the CMB (Das et al., 2011; Sherwin et al., 2011) and the first measurement of the bulk motion of massive clusters of galaxies using the kinematic Sunyaev-Zeldovich effect (Hand et al., 2012). The latest public data release from the collaboration in 2020, Data Release 4 (DR4; see Fig. 2),¹ has generated great interest in the cosmology research community and received widespread coverage in the international press (e.g., Mortillaro, 2020; Amos, 2020). The maps and statistical data products released as part of DR4 yielded excellent agreement with the standard model of cosmology and provided the first independent check by a high resolution experiment of the Hubble constant with CMB data (Aiola et al., 2020; Choi et al., 2020). Numerous ACT papers based on DR4 and preliminary versions of the next data release, DR5, have appeared or are in preparation.²

ACT and other CMB experiments have not only benefited from Canadian DRI but have spurred its development: one of the reasons SciNet was created was to provide the computational resources required by these collaborations. Today, a new opportunity for DRI innovation in Canada is presented by the Simons Observatory (SO), an upcoming CMB experiment. For its initial five-year survey (2024–29), SO will consist of one 6 m large aperture telescope (LAT) and three 50 cm small aperture telescopes (SAT; Lee et al., 2019). Together they will deliver an order-of-magnitude better sensitivity than current CMB experiments, with the lower-resolution SATs geared towards detecting the signature of primordial gravity waves that would be decisive evidence that the Universe underwent a rapid period of inflation in the first instants after the big bang, while the higher-resolution LAT will measure the total mass of the three neutrino species, learn about the nature of dark matter and dark energy, discover many thousands of clusters of galaxies, and much more (Ade et al., 2019). The increase in sensitivity provided by SO is enabled by fielding more than ten times the number of detectors in ACT, but this advance comes with a commensurate increase in the volume and complexity of the data that need to be processed.

We propose that the New Digital Research Infrastructure Organization (NDRIO) help enable SO science by

¹Available at https://lambda.gsfc.nasa.gov/product/act/actpol_prod_table.cfm.

²See <https://act.princeton.edu/publications>.




Figure 2: A section of the map of E-mode polarisation of the CMB from ACT’s DR4 (Aiola et al., 2020; Choi et al., 2020). This map was made with SciNet’s Niagara cluster. The width shown here is 25 degrees, corresponding to 20 billion light years. The colours represent small variations in the amount of polarisation ($\pm 25 \mu\text{K}$) and the size of these features was used to make a new measurement of the Hubble constant. (Image credit: ACT Collaboration.)

providing support to make CMB maps from the LAT data. This would build on Canada’s world-class leadership in developing CMB experiments and processing and analysing their data. The DRI required by SO represents a significant increase over current needs, and would be an impressive contribution to the collaboration whose data products will be essential to cosmology research for the coming decade. After a review of the current use of Canadian DRI for ACT data processing and the particular challenges we are beginning to face (§2), we present the DRI required for the reduction of SO LAT data (§3). Finally, we make concrete recommendations for how advanced research computing (ARC) and research data management (RDA) can be integrated to meet the needs of a large-scale project like SO (§4).

2 Current Use of DRI and Associated Challenges

DRI Resources Currently Required: In order to contextualise the needs of SO, we provide here a summary of ACT’s anticipated use of resources on SciNet’s Niagara cluster for the coming year. By far the majority of cycles consist in reducing raw data to maps of the CMB sky in temperature and polarisation (see Table 1). ACT’s $\sim 3,000$ detectors produce about 15 terasamples per year, for a total of 130 TB for the 2017–20 dataset being processed. The map-making software, largely written by ACT researchers, uses efficient parallel algorithms, particularly pre-conditioned conjugant gradient (PCG) solvers, to complete the enormous matrix operations required to reconstruct the maps while optimally removing the noise from the celestial signal. The non-trivial challenges associated with doing this over a large fraction of the sky (40%) at \sim arcminute resolution have made heavy use of the highly reliable MPI communication provided by Niagara. Map-making of ACT’s 2017–20 observations will take about 1280 core years,³ with about 20 TB of scratch space needed for storing intermediate data products during the process; the final products are ~ 5 TB. Memory needs are relatively modest and Niagara’s ~ 4 GB RAM per core are quite sufficient. In addition to map-making, Niagara is used to determine the telescope pointing solution using time-series analyses of bright point sources (~ 120 core years). Furthermore, interpretation of the final products, especially the gravitational lensing of the CMB, relies on processing of simulated data, requiring about ~ 250 core-years. To date, all of the foregoing work has been done by about 30 researchers under a CC allocation led by principal investigator (PI) J. R. Bond. These researchers, who range from students to senior faculty, are based at both Canadian and international institutions. The same allocation, currently being used near capacity, supports other important CMB projects. The SPIDER collaboration (~ 20 researchers) uses it to do similar map-making to ACT on larger angular scales. In the coming year they aim to finalise maps from a 16-day flight of their balloon-borne telescope in 2015 (Gualtieri et al., 2018); together with generating the requisite simulations, this will require about 180 core-years with ~ 13 TB

³Here and throughout we refer to CPU, rather than GPU, processing.

Project	ACT			SPIDER	Theory Sims	Total
	<i>map-making</i>	<i>pointing</i>	<i>lensing analysis</i>			
Core Years	1280	120	250	180	250	2080
Storage (TB)	130			13	12	155
No. Researchers	~30			~20	~15	~65

Table 1: Anticipated use of DRI resources for CMB research on SciNet’s Niagara cluster in 2021–22.

of storage space. Finally, about 15 researchers employ the allocation for creating sophisticated models of secondary components in the CMB, such as clusters of galaxies, point sources and the effects of possible non-Gaussian physics in the early Universe. These models inform not only analysis of the CMB but also other cosmological datasets, such as those from the Canadian Hydrogen Intensity Mapping Experiment (Bandura et al., 2014) or the *Euclid* satellite (Racca et al., 2016), which have significant Canadian membership. In the coming year this work will take about 250 core years and 12 TB of storage.

Current Limitations and Challenges: Thus far, the CC allocation model has served the CMB community extremely well. However, two limitations of the current system are beginning to emerge. First, CC allocations are tied to single PIs, which can be unwieldy for large collaborations—or, in the scenario described above, for more than one large collaboration sharing an allocation. The allocation PI is not necessarily the PI of the collaboration; if, for some reason, the former must step down, it is not clear how the collaboration’s work can move forward without the significant interruption of having a new PI apply for a new allocation. Furthermore, within just a single collaboration like ACT there are several distinct computational tasks (c.f., Table 1) that are nonetheless tightly coupled. Currently there are two possibilities: keep them under a single allocation, where they may compete for cycles, or put them under separate allocations, where it becomes more difficult to integrate their results. Neither is fully satisfactory. Second, the one-year allocation length is short compared with the years of continual effort now required to process a single CMB dataset. Even though allocation renewals can be ‘fast tracked’ for two more years after an initial successful application, the cadence of allocation cycles is becoming incommensurate with experiment timelines. To address these limitations, we propose the introduction of a ‘team allocation’, described in §4, below.

3 Looking to the Future: Computing Requirements for the SO LAT

Building on the Legacy of ACT: The SO LAT is similar to ACT in terms of sky coverage (~40%), angular resolution and frequencies observed. Doing SO map-making on Canadian DRI thus makes eminent practical sense, since much of the research software (RS) will be similar and build on the demonstrated success of ACT computation. The LAT will do a five years survey using about ten times more detectors than ACT, and this increased data volume will be the main challenge for DRI. It requires a major step forward in terms of storage, computing cycles and total amount of resources used per single job. Let us now quantify these requirements.

Required DRI for the SO LAT: The SO LAT will produce roughly 1.5 PB of raw time-ordered data over its five year survey.⁴ Each year of data (called a ‘season’) will be continuously analysed for at least two subsequent years; indeed it is conceivable that all 1.5 PB will need to remain ‘spinning’—or if on tape, seamlessly transferred to disk upon read—for a seven year period during which the full dataset will be reprocessed with improvements to the reduction pipelines.⁵ We base this upon experience from ACT that as sensitivity grows with increased data volume, new systematic effects emerge that need to be understood, modelled and accounted for with sophisticated upgrades to the software pipelines. In the high-precision realm of SO science, this process will be essential. Final science-quality products and associated simulations will also require storage at the level of a few hundred TB and

⁴This figure is less than one gets from scaling the ACT volume by a 10× increase in detectors, due to lower sampling rates on SO.

⁵The number seven comes from the full period of official SO operations: five years of observations and data reduction, and two extra years to complete the final delivery of the products.

will need to be made accessible to a large number of collaborators and to the broader community. In addition to the large amount of storage, a significant increase in computational cycles is required. From estimates based on the performance of the ACT code, we expect a computational cost of 300 core years to reduce a single season of LAT data. Additional RAM (on the order of 10's of TB total) can reduce this somewhat, both by eliminating repeated calculations and enabling more algorithms. Nonetheless, the CPU usage number only represents the final pass at the data: before production runs, there is a critical phase of data exploration and characterisation of systematics, as mentioned above. Even though only a subset of the yearly data is processed, it must be done multiple times. Our rough estimate is that such exploratory work amounts to about 20 times the cycles of a full season run. This brings the total to about 6,000 core-years per season of data. In the scenario of a cumulative analysis of the dataset, the annual requirement will have increased to roughly 30,000 core-years by 2029, or 18 times higher than the anticipated cycles needed for ACT processing in 2021–22 (see Table 1).

Coordination with International DRI: SO is a large international collaboration, some of whose members have access to non-Canadian DRI, or are in positions to make strong bids in competitive applications for their use. Based on previous experience, the best approach is not to make use of a single resource but to strategically coordinate among them: for instance, while almost all the principal map-making for ACT has been done on Niagara, other computing has been done on clusters at the U.S. National Energy Research Scientific Computing Center (NERSC), at the university of KwaZulu-Natal and at Princeton University. This multi-pronged approach has proven extremely effective for avoiding single-point failures and for distributing diverse types of computation among different DRI. For SO, we envision a scenario where it is not Canadian *or* international DRI, but Canadian DRI *with* international DRI. As a concrete example, we might envision running time-domain simulations of our raw data and map-making of the lower resolution SAT data at NERSC, while using Canadian DRI for map-making of the LAT data as proposed above. The LAT and SATs are sufficiently decoupled that their data can be processed almost completely independently of each other, making such a division of labour trivial. Nonetheless, we should add that the coordination of DRI outlined in this White Paper (i.e., LAT map-making in Canada and other processing at NERSC) is not the only possibility. There are other scenarios in which SAT map-making and other computational tasks are done with NDRIO DRI that would utilise similar resources to those we outline in this White Paper. Our intent has been to focus on a well-motivated plan that provides a clear indication of DRI requirements for upcoming CMB processing.

4 Bridging the Gap: Recommendations for NDRIO

Let us now make some specific recommendations for NDRIO that would enable us to capitalise on the singular opportunity to process SO LAT data with Canadian DRI. An obvious *sine qua non* would be securing an allocation on ARC that meets the requirements for storage and computing power outlined above (§3), namely:

- capacity for 30,000 core years of processing (baselined to Niagara's CPUs),
- 1.5 PB of spinning storage during processing,
- 1.5 PB of cold storage for data recovery and
- several 100 TB of storage for final, publicly available data products (maps and associated simulations)

Additionally, we propose the following features for NDRIO ARC and RDM:

1. Team Allocations: To overcome the limitation of the single-PI allocation model described in §2, we recommend the addition of a *team allocation* with a 7 year duration. The allocation would be targeted at SO-sized collaborations, and would be applied for and managed by a small group of Canada-based PIs, or team leaders, from within the collaboration. A team allocation would be used for the largest and longest-term computing tasks of the collaboration—in the case of SO, this would consist of map-making. Single PI allocations would still exist, and members of the collaboration would be encouraged to secure their own allocations for short-term, ancillary projects requiring HPC. The length of single-PI allocations need not be as long as team allocations.

2. Collocation of Public Data Hosting and ARC: In the hybrid allocation model described above, there would have to be RDM capabilities such that members of the collaboration with single-PI allocations could seamlessly access, add to or update data products produced with the team allocation. Most basically this could simply occur through intelligent management of file system permissions, though one might imagine RDM solutions that aid sharing of various stages of data products within a team allocation and across authorised single-PI allocations. However, we propose a much more ambitious vision of RDM that enables collaborations to release final, public data products to the broader research community. Given the enormous size of these products (hundreds of TB for SO: see §3), this would mean more than just hosting files for download as per the current model; though this would still be important, it could be impracticable when the whole dataset is involved. Ideally, anyone with a single-PI or team allocation on NDRIO ARC could directly access these products—in other words, they would be collocated with the ARC. One could even envision RDM that automatically mirrors public data products at multiple HPC facilities across the country for the broadest possible access; furthermore, we recommend that there be a pathway for non-Canadian researchers to apply for NDRIO ARC resources in order to amplify the international impact of the data products by maximising their use.

5 Conclusions

For years now, cosmologists have turned to Canadian DRI for the critical work of processing raw time streams from CMB experiments into the science-grade data products that have advanced our knowledge of the Universe. With the advent of next-generation CMB experiments in the coming years, we have an exciting opportunity to build upon this legacy and adapt to the new computational demands that larger observatories present. To illustrate this we have focused on the transition from ACT to the SO LAT, and presented a scenario in which a significant contribution to upcoming CMB science can be achieved with future Canadian DRI within an international context. As mentioned in the Introduction, we believe the recommendations outlined in this paper would benefit not only the SO collaboration, but other future CMB experiments, and provide data products that will enhance our understanding of cosmology for many years to come.

Acknowledgements We thank Michael Nolta for helpful advice and feedback during preparation of this paper.

References

- Ade, P., Aguirre, J., Ahmed, Z., et al. 2019, *J. Cosmology Astropart. Phys.*, 2019, 056, doi: [10.1088/1475-7516/2019/02/056](https://doi.org/10.1088/1475-7516/2019/02/056)
- Aiola, S., Calabrese, E., Maurin, L., et al. 2020, arXiv e-prints, arXiv:2007.07288. <https://arxiv.org/abs/2007.07288>
- Amos, J. 2020, BBC News. <https://www.bbc.com/news/science-environment-53420433>
- Bandura, K., Addison, G. E., Amiri, M., et al. 2014, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 9145, *Ground-based and Airborne Telescopes V*, ed. L. M. Stepp, R. Gilmozzi, & H. J. Hall, 914522, doi: [10.1117/12.2054950](https://doi.org/10.1117/12.2054950)
- Barmby, P., Gaensler, B., Dobbs, M., et al. 2020, Final report of the Canadian Astronomical Society’s 2020 Long Range Plan for Canadian Astronomy (LRP2020), Tech. rep., Canadian Astronomical Society. https://casca.ca/wp-content/uploads/2020/12/LRP2020_December2020-1.pdf
- Choi, S. K., Hasselfield, M., Ho, S.-P. P., et al. 2020, arXiv e-prints, arXiv:2007.07289. <https://arxiv.org/abs/2007.07289>
- Das, S., Sherwin, B. D., Aguirre, P., et al. 2011, *Phys. Rev. Lett.*, 107, 021301, doi: [10.1103/PhysRevLett.107.021301](https://doi.org/10.1103/PhysRevLett.107.021301)
- Gualtieri, R., Filippini, J. P., Ade, P. A. R., et al. 2018, *Journal of Low Temperature Physics*, 193, 1112, doi: [10.1007/s10909-018-2078-x](https://doi.org/10.1007/s10909-018-2078-x)
- Hand, N., Addison, G. E., Aubourg, E., et al. 2012, *Phys. Rev. Lett.*, 109, 041101, doi: [10.1103/PhysRevLett.109.041101](https://doi.org/10.1103/PhysRevLett.109.041101)
- Henderson, S. W., Allison, R., Austermann, J., et al. 2016, *Journal of Low Temperature Physics*, 184, 772, doi: [10.1007/s10909-016-1575-z](https://doi.org/10.1007/s10909-016-1575-z)
- Hlozek, R., Bond, J. R., Chapman, S., et al. 2019, in *Canadian Long Range Plan for Astronomy and Astrophysics White Papers*, Vol. 2020, 50, doi: [10.5281/zenodo.3825611](https://doi.org/10.5281/zenodo.3825611)
- Lee, A., Abitbol, M. H., Adachi, S., et al. 2019, in *Bulletin of the American Astronomical Society*, Vol. 51, 147. <https://arxiv.org/abs/1907.08284>
- Mortillaro, N. 2020, CBC News. <https://www.cbc.ca/news/technology/age-universe-science-cmb-1.5653450>
- Racca, G. D., Laureijs, R., Stagnaro, L., et al. 2016, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 9904, *Space Telescopes and Instrumentation 2016: Optical, Infrared, and Millimeter Wave*, ed. H. A. MacEwen, G. G. Fazio, M. Lystrup, N. Batalha, N. Siegler, & E. C. Tong, 99040O, doi: [10.1117/12.2230762](https://doi.org/10.1117/12.2230762)
- Sherwin, B. D., Dunkley, J., Das, S., et al. 2011, *Phys. Rev. Lett.*, 107, 021302, doi: [10.1103/PhysRevLett.107.021302](https://doi.org/10.1103/PhysRevLett.107.021302)

Swetz, D. S., Ade, P. A. R., Amiri, M., et al. 2011, ApJS, 194, 41, doi: 10.1088/0067-0049/194/2/41
Thornton, R. J., Ade, P. A. R., Aiola, S., et al. 2016, ApJS, 227, 21, doi: 10.3847/1538-4365/227/2/21